Communication Cost for Updating Linear Functions when Message Updates are Sparse: Connections to Maximally Recoverable Codes

N. Prakash and Muriel Médard

arXiv:1605.01105v2 [cs.IT] 8 Jun 2016

Abstract

We consider a communication problem in which an update of the source message needs to be conveyed to one or more distant receivers that are interested in maintaining specific linear functions of the source message. The setting is one in which the updates are sparse in nature, and where neither the source nor the receiver(s) is aware of the exact *difference vector*, but only know the amount of sparsity that is present in the difference-vector. Under this setting, we are interested in devising linear encoding and decoding schemes that minimize the communication cost involved. We show that the optimal solution to this problem is closely related to the notion of maximally recoverable codes (MRCs), which were originally introduced in the context of coding for storage systems. In the context of storage, MRCs guarantee optimal erasure protection when the system is partially constrained to have local parity relations among the storage nodes. In our problem, we show that optimal solutions exist if and only if MRCs of certain kind (identified by the desired linear functions) exist. We consider point-to-point and broadcast versions of the problem, and identify connections to MRCs under both these settings.

I. Introduction

We consider a communication problem in which an update of the source message needs to be communicated to one or more distant receivers that are interested in maintaining specific functions of the actual source message. The setting is one in which the updates are sparse in nature, and where neither the source nor receivers are aware of the *difference-vector*. Under this setting, we are interested in devising encoding and decoding schemes that allow the receiver to update itself to the desired function of the new source message, such that the communication cost is minimized.

A. Pont-to-Point Setting

The system for the point-to-point case is shown in Fig. 1. Here, the n-length column-vector $X^n \in \mathbb{F}_q^n$ denotes the initial source message, where $\mathbb{F}_q$ denotes the finite field of q elements. The receiver maintains the linear function AX^n, where A is an $m \times n$ matrix, $m \leq n$, over $\mathbb{F}_q$. Let the updated source message be given by $X^n + E^n$, where E^n denotes the difference-vector, and is such that Hamming wt.$(E^n) \leq \epsilon$, $0 \leq \epsilon \leq n$. We say that the vector E^n is ϵ-sparse. We consider linear encoding at the source using the $\ell \times n$ matrix H. We assume that the source is aware of the function A and the parameter ϵ, but does not know the vector E^n. The goal is to encode $X^n + E^n$ at the source such that the receiver can update itself to $A(X^n + E^n)$, given the source encoding and AX^n. Assuming that the parameters n, q, ϵ and A are fixed, we define the communication cost as the parameter ℓ, which is the number of symbols generated by the linear encoder H. The goal is to design the encoder and decoder so to minimize ℓ. We assume a zero-probability-of-error, worst-case-scenario model; in orther words, we are interested in schemes which allow error-free decoding for every $X^n \in \mathbb{F}_q^n$ and ϵ-sparse $E^n \in \mathbb{F}_q^n$.

The problem is in part motivated by the setting of distributed storage systems (DSSs) that use linear erasure codes for data storage, and where the underlying uncoded file is subject to frequent but tiny updates. In a scenario where multiple users share the DSS, it is often the case that a user applies his updates on a version of the file

N. Prakash and Muriel Médard are with the Research Laboratory of Electronics, Massachusetts Institute of Technology, USA (email: {prakashn, medard}@mit.edu). The results in this article were presented in part as an invited talk at 53rd Annual Allerton Conference on Communication, Control, and Computing, Sept 29-Oct 2, 2015, Allerton Park and Retreat Center, Monticello, IL, USA. The work is in part supported by AFOSR under grant number FA9550-14-1-043.

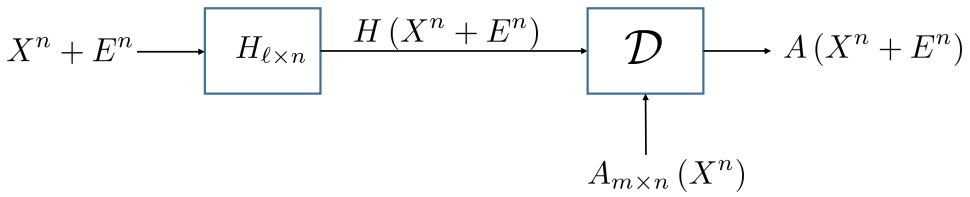

Fig. 1. The system model for function update in the point-to-point setting.

which was obtained locally from a different user, and which is not present in the DSS. In our model, we consider updates as substitutions in the file, which is a reasonable assumption in scenarios when the file-size is much larger than the quantum of updates. With respect to Fig. 1, X^n represents the version of the file that is stored in the DSS. The user has access to a locally obtained version $X^n + E_1^n$ that is not stored in the DSS, and updates it to $X^n + E_1^n + E_2^n = X^n + E^n$, where $E^n = E_1^n + E_2^n$ represents the overall update that the file has undergone. Further, our interest is in DSSs like peer-to-peer DSS where the user typically establishes multiple connections with various nodes in order to update their coded data. Figure 1 illustrates a scenario in which the user updates the coded data corresponding to one of the storage nodes that uses matrix A as its erasure coding coefficients. We illustrate the usage of the matrix A via the following example.

Example 1: Consider a DSS that uses an $[N, K]$ linear code for storing data across N storage nodes. The data file X^n is striped before encoding, the various stripes are individually encoded and stacked against each other to get the overall coded file. Let $\mathbf{a} = [a_1 \ a_2 \ \cdots a_K], a_i \in \mathbb{F}_q$ denote the coding coefficients corresponding to one of the storage nodes. Assuming that the first K symbols of the vector X^n correspond to the first stripe, symbols $K+1$ to $2K$ correspond to the second stripe and so on, the overall coding matrix A (taking into account all stripes) is given by

$$A = \begin{bmatrix} \mathbf{a} & & & \\ & \mathbf{a} & & \\ & & \ddots & \\ & & & \mathbf{a} \end{bmatrix}. \tag{1}$$

Here we assume that $K|n$, and thus there are $m = \frac{n}{K}$ stripes in the file. The parameter m also represents the number of coded symbols stored by any one of the N nodes. The matrix A is of size $m \times mK$. The scenario of interest is one in which the number of coded symbols m is much larger than the sparsity parameter ϵ, which is the number of symbols that get updated in the uncoded file. The goal is to design the encoder H and the decoder $\mathcal{D}$ that minimizes the communication cost.

B. Broadcast Setting

The system model naturally extends to broadcast settings consisting of a single source and multiple destinations that are interested in possibly different functions of the source message. In our work, we study the broadcast setting for the case of two destinations (see Fig. 2), where the receivers initially hold functions AX^n and BX^n, and get updated to $A(X^n + E^n)$ and $B(X^n + E^n)$, respectively. The broadcast setting is motivated by settings in which two of the storage nodes in a peer-to-peer DSS are nearly co-located, as far as a (distant) user is concerned. In this case, the question of interest is to identify whether it is possible to send one encoded packet $H(X^n + E^n)$ to simultaneously update the coded data of both the storage nodes. We will next present two examples of the broadcast setting. As we will see later in this document, the first example is an instance in which broadcasting does not help to reduce communication cost, i.e., it's optimal to individually update the two destinations. The second example is an instance in which it is beneficial to broadcast than individually update the two destinations.

Example 2: As in Example 1, we consider striping and encoding of a data file by an $[N, K]$ linear code over $\mathbb{F}_q$ with $K \geq 2$. We apply the broadcast model to simultaneously update the contents of 2 out of the N storage nodes. Let us assume that the K-length coding vectors associated with two of the storage nodes are given by

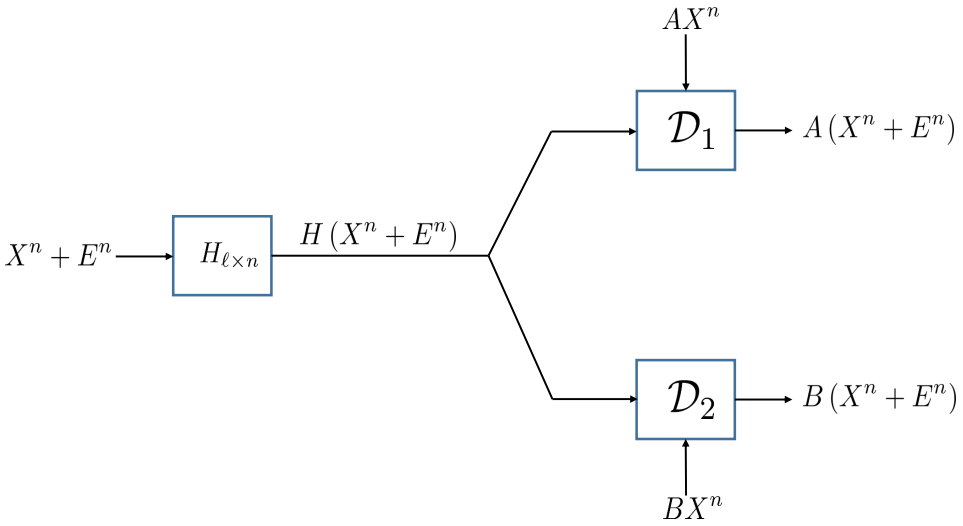

Fig. 2. The system model for function updates in the broadcast setting involving two receivers.

$\mathbf{a} = [a_1\ a_2\ \ldots a_K], \mathbf{b} = [b_1\ b_2\ \ldots b_K], a_i, b_i \in \mathbb{F}_q$. The overall coding matrices A and B corresponding to the m stripes are then given by

$$A = \begin{bmatrix} \mathbf{a} & & & \\ & \mathbf{a} & & \\ & & \ddots & \\ & & & \mathbf{a} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} \mathbf{b} & & & \\ & \mathbf{b} & & \\ & & \ddots & \\ & & & \mathbf{b} \end{bmatrix}, \tag{2}$$

where both A and B have size $m \times mK$. We will see later that for this example, in order to achieve the optimal communication cost, the source must transmit as though it is individually updating the two storage nodes, i.e., there is no added benefit due to broadcasting.

Example 3: In this example, we consider striping and encoding of a data file by a *minimum bandwidth regenerating* (MBR) code [1] that is obtained via the product-matrix construction described in [2]. Regenerating codes [1] are codes specifically designed for data storage, and they allow the system designer to trade-off storage overhead against *repair bandwidth* for a given level of fault tolerance. MBR codes have the best possible repair bandwidth (at the expense of storage overhead) for a given level of fault tolerance. We first give a quick introduction to a specific instance of an MBR code in [2], and use that in our broadcast setting.

The MBR code: Consider an $[N = 5, K = 3, D = 4](\alpha = 4, \beta = 1)$ MBR code over $\mathbb{F}_q$ that encodes 9 symbols into 20 symbols and stores across $N = 5$ nodes such that each node holds $\alpha = 4$ symbols. The code has the property that the contents of any $K = 3$ nodes are sufficient to reconstruct the 9 uncoded symbols. Also, the contents of any one of the $N = 5$ node can be recovered by connecting to any set of $D = 4$ other nodes[1] and downloading $\beta = 1$ symbols from each of them. The latter property is known as the node-repair property of the regenerating code. Let $[m_1, \ldots, m_9]$ denote the vector of 9 message symbols, where $m_i \in \mathbb{F}_q$. The encoding is described by a product-matrx construction as follows:

$$C_{N\times\alpha} = \Psi_{N\times D} M_{D\times\alpha} \tag{3}$$

$$= \begin{bmatrix} \underline{\psi}_1 \\ \underline{\psi}_2 \\ \underline{\psi}_3 \\ \underline{\psi}_4 \\ \underline{\psi}_5 \end{bmatrix} \begin{bmatrix} m_1 & m_2 & m_3 & m_7 \\ m_2 & m_4 & m_5 & m8 \\ m_3 & m_5 & m_6 & m_9 \\ m_7 & m_8 & m_9 & 0 \end{bmatrix}. \tag{4}$$

[1]In our example, one connects to all the remaining $N - 1 = 4$ nodes, since $D = 4$.

Here, M denotes the $D \times \alpha$ message matrix whose entries are populated from among the 9 message symbols, in a specific manner, as shown in (4). Note that the matrix M is symmetric. The $N \times D$ encoding matrix Ψ can be chosen as a Vandermonde matrix under the product-matrix framework. The vector $\underline{\psi}_i, 1 \leq i \leq 5$ denotes the i^{th} row of Ψ. In this example, we assume that Ψ is given as follows:

$$\Psi \;=\; \begin{bmatrix} 1 & \gamma & \gamma^2 & \gamma^3 \\ 1 & \gamma^2 & \gamma^4 & \gamma^6 \\ 1 & \gamma^3 & \gamma^6 & \gamma^9 \\ 1 & \gamma^4 & \gamma^8 & \gamma^{12} \\ 1 & \gamma^5 & \gamma^{10} & \gamma^{15} \end{bmatrix}, \tag{5}$$

where we pick γ as a primitive element in $\mathbb{F}_q$. The $N \times \alpha$ matrix C is the codeword matrix, with the i^{th} row representing the contents that get stored in the i^{th} node, $1 \leq i \leq 5$.

Striping using the MBR code, and the associated coding matrices A and B: As before, we let X^n to denote the uncoded data file, which is divided into m stripes each of length 9 symbols. Also, recall that in our model of striping, the first stripe (for the current example) consists of the first 9 symbols of X^n, the second stripe consists of symbols $X_{10}, \ldots, X_{18}$, and so on. The length n is given by $n = 9m$. Consider the case where we use the broadcast model to update the contents of the first two nodes. The overall coding matrices A and B, both of size $4m \times 9m$, corresponding to the first and the second node are respectively given by

$$A \;=\; \begin{bmatrix} A_1 & & & \\ & A_1 & & \\ & & \ddots & \\ & & & A_1 \end{bmatrix} \quad \text{and} \quad B \;=\; \begin{bmatrix} B_1 & & & \\ & B_1 & & \\ & & \ddots & \\ & & & B_1 \end{bmatrix}, \tag{6}$$

where

$$A_1 \;=\; \begin{bmatrix} 1 & \gamma & \gamma^2 & & & & \gamma^3 & & \\ & 1 & & \gamma & \gamma^2 & & & \gamma^3 & \\ & & 1 & & \gamma & \gamma^2 & & & \gamma^3 \\ & & & & & & 1 & \gamma & \gamma^2 \end{bmatrix} \quad \text{and} \quad B_1 \;=\; \begin{bmatrix} 1 & \gamma^2 & \gamma^4 & & & & \gamma^6 & & \\ & 1 & & \gamma^2 & \gamma^4 & & & \gamma^6 & \\ & & 1 & & \gamma^2 & \gamma^4 & & & \gamma^6 \\ & & & & & & 1 & \gamma^2 & \gamma^4 \end{bmatrix}. \tag{7}$$

As we shall later, such broadcasting, rather than individually transmitting to the two destinations, can reduce the total communication cost for this example by nearly 12%.

C. Connection to Maximally Recoverable Codes

In this paper, we identify necessary and sufficient conditions for solving the function update problems in Fig. 1, Fig. 2, for general matrices A and B. We show that the existence of optimal solutions under both these settings is closely related to the notion of *maximally recoverable codes*, which were originally studied in the context of coding for storage systems [3], [4], [5]. In the context of storage, MRCs form a subclass of a broader class of codes known as locally repairable codes (LRC) [6], [7]. An $[n, k]$ linear code $\mathcal{C}$ is called an LRC with locality r if each of the n code symbols is recoverable as a linear combination of at most r other symbols of the code. Qualitatively, an LRC is said to be maximally recoverable (a formal definition appears in Section II) if it offers optimal *beyond-minimum-distance* correction capability. When restricted to the setting of LRCs, MRCs are known as partial-MDS codes [4] or maximally recoverable codes with locality [5]. MRCs with locality are used in practical DSSs like Windows Azure [8].

The usage of the concept of MRCs in our work is more general, and not necessarily restricted to the context of LRCs. Below, we define the notion of *maximally recoverable subcode* of a given code, say $\mathcal{C}_0$.

Definition 1 (Maximally Recoverable Subcode $\mathcal{C}$ of $\mathcal{C}_0$): Let $\mathcal{C}_0$ denote an $[n, t]$ linear code over $\mathbb{F}_q$ having a generator matrix G_0. Also, consider an $[n, k]$ subcode $\mathcal{C}$ of $\mathcal{C}_0$ for some $k \leq t$, and let G denote a generator matrix of $\mathcal{C}$. The code $\mathcal{C}$ will be referred to as a maximally recoverable subcode (MRSC) of $\mathcal{C}_0$ if for any set $S \subseteq [n], |S| = k$ such that rank $(G_0|_S) = k$, we have rank $(G|_S) = k$.

Example 4: Consider an $[n = 9, t = 3]$ binary code $\mathcal{C}_0$ whose generator check matrix G_0 is given by

$$G_0 \;=\; \begin{bmatrix} 1 & 1 & 1 & & & & & & \\ & & & 1 & 1 & 1 & & & \\ & & & & & & 1 & 1 & 1 \end{bmatrix}. \tag{8}$$

Next, consider the $[n = 9, k = 2]$ subcode $\mathcal{C}$ of $\mathcal{C}_0$ having a generator matrix G given by

$$G \;=\; \begin{bmatrix} 1 & 1 & 1 & & & & 1 & 1 & 1 \\ & & & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}. \tag{9}$$

It is straightforward to verify that $\mathcal{C}$ is an MRSC of $\mathcal{C}_0$.

D. Summary of Results

Following is a summary of the results obtained in this paper.

1) For the point-to-point setting in Fig. 1, the communication cost is lower bounded by[2]

$$\ell \;\geq\; \max(2\epsilon, \text{rank}(A)). \tag{10}$$

If C_A and C_H denote n-length linear codes respectively generated by the rows of the matrices A and H, then under optimality (i.e., achieving equality in (10)), we show that the code C_H must necessarily be a MRSC of C_A. An achievable scheme based on MRSCs is presented to establish the optimality of the bound in (10). For a general matrix A, our achievability is guaranteed only when the field size q used in the model (see Fig. 1) is sufficiently large.

2) An explicit low field size encoding matrix H is provided for the setting considered in Example 1. Recall that in Example 1, we use an arbitrary $[N, K]$ linear code to stripe and store the data. The MRSC C_H in this case corresponds to MRCs with locality, where the local codes are *scaled repetition codes*.

3) We identify necessary and sufficient conditions for solving the broadcast setting given in Fig. 2. Let C_A and C_B respectively denote the linear block codes generated by the rows of the matrices A and B. For the special case when the codes C_A and C_B intersect trivially, there is no benefit from broadcasting, i.e., the encoder must transmit as though it is transmitting individually to the two receivers, and thus the optimal communication is given by

$$\ell \;\geq\; \max(2\epsilon, \text{rank}(A)) + \max(2\epsilon, \text{rank}(B)). \tag{11}$$

For the general case when C_A and C_B have a non-trivial intersection, the optimal communication cost can be less than what is given by (11). The expression for the optimal communication cost appears in Theorem 5.5. Like in the point-to-point case, optimality is guaranteed only when the field size q is sufficiently large.

4) Our achievability scheme for the general case of the broadcast setting involves finding answer to the following sub-problem : Given an $[n, t]$ code $\mathcal{C}_0$ and an $[n, s]$ subcode $\widehat{\mathcal{C}}$ of $\mathcal{C}_0$, can we find an $[n, k]$ MRSC $\mathcal{C}$ of $\mathcal{C}_0$ such that $\widehat{\mathcal{C}}$ is a subcode of $\mathcal{C}$. We refer to this as the problem of finding *sandwiched MRSCs*. Here the parameters s, k, t satisfy the relation $s \leq k \leq t$. We present two different techniques that show the existence of sandwiched MRSCs (under certain trivial necessary conditions on the code $\mathcal{C}_0$). The first technique is a parity-check matrix based approach, and relies on the existence of non-zero evaluations of a multivariate polynomial over a large enough finite field. The second technique is a generator matrix based approach, uses the concept of *linearized polynomials*, and yields an explicit construction.

E. Related Work

Coding for File Synchronization : In [9], the authors consider the problem of oblivious file synchronization, under the substitution model, in a distributed storage setting employing linear codes. The model is one in which one of the coded storage nodes gets updated with the help of other coded storage nodes, who have already undergone updates. The authors considers the problem of jointly designing the storage code, and also the update scheme. The updates, like in our setting, is carried out under the assumption that only the sparsity of updates is known, and

[2]The bound in (10) was already proved in [9] for an average case setting.

not the actual updates themselves. Optimal solutions are presented when the storage code is assumed to an MDS code. A recent version of this work [10] extends the work to the setting of regenerating codes, with the restriction that only one symbol ($\epsilon = 1$) is obliviously updated. The key difference between these works and ours is that we consider designing optimal update schemes for an arbitrary linear function, while the above works assume a specific structure of the storage code (e.g., MDS in [9]). Please also see Remark 1 for a comparison of the converse statements appearing in the two papers.

A second work which is related to ours appears in [11], where the authors consider the problem of designing protocols for simultaneously optimizing storage performance as well as communication cost for updates. Contrary to our setting, the updates are modeled as insertions/deletions in the file. A similarity with our work (like the work in [9]) is that the destination node holds a coded form of the data (which can be considered as linear function of the uncoded data), and is interested in updating the coded data. The goal is to communicate and update the coded file in such a way that reconstruction/repair properties in the storage system are preserved. Their protocol permits modifications to the structure of the storage code itself for optimizing the communication cost. Note that in our model we optimize the communication cost under the assumption that the destination is interested in the *same* function A of the updated data as well.

The problem of synchronizing files under the insertion/deletion model has been previously studied, using information-theoretic techniques, in [12], [13]. Recall that our model uses substitutions, rather than insertions/deletions, for specifying updates. In [12], a point-to-point setting is considered in which the decoder has access to a partially-deleted version of the input sequence, as side information. The authors calculate the minimum rate at which the source needs to compress the (undeleted) input sequence so as to permit lossless recovery (i.e., vanishing probability of error for large block-lengths) at the destination. The model of [13] allows both insertions and deletions in the original file to get the updated file. Further, both the source and the destination have access to the original file. Under this setting, the authors study the minimum rate at which the updated file must be encoded at the source, so that lossless recovery is possible at the destination. The problem of synchronizing edited sequences is also considered in [14], where the authors allow insertions, deletions or substitutions, and assume a zero-error recovery model. An important difference between our work and all the works mentioned above is that while we are interested in a specific linear function of the input sequence, all of the above works assume recovery of the entire source sequence at the destination.

Maximally Recoverable Codes: The notion of maximal recoverability in linear codes was originally introduced in [3]. Low field size constructions of partial MDS codes (another name for MRCs with locality) for specific parameter sets codes appear in [4], [15], [16]. In the terminology of partial MDS codes, these three works respectively provide low field size constructions for the settings up to one, two and three global parities, and any number of local parities. Explicit constructions based on linearized polynomials for the case of one local parity and any number of global parities appear in [5]. Identifying low-field-size constructions of partial MDS codes for general parameter sets remain an open problem. Various approximations of partial MDS codes like Sector Disk codes [4], [15], [16], STAIR codes [17], and partial-maximally-recoverable-codes [18] have been considered in literature, and these permit low field size constructions for larger parameter sets. Other known results on maximally recoverable codes include expressions for the weight enumerators of MRCs with locality [19], and the generalized Hamming weights (GHWs) of the MRSC $\mathcal{C}$ in terms of the GHWs of the code $\mathcal{C}_0$ [20].

Zero-Error Function Computation : Finally, we note that the problem considered in this paper can be considered as one of zero-error function computation. Even though the existing literature does not directly address the problem that we study here, we review some of the relevant works on zero-error function computation, in point-to-point as well as network settings. In [21], the authors examine the problem of computing a general function with zero-error, under worst-case as well as average-case models in the context of sensor networks. One of the problems considered there involves a point-to-point setting, where a source having access to x needs to encode and transmit to a destination which has access to side information y. The destination is interested in computing the function $f(x, y)$ with zero-error in the worst-case model. Though the point-to-point model in Fig. 1 can be considered as a special case of the setting in [21], our assumptions regarding the specific nature of the decoder side-information, and the use of linear encoding enable us to show deeper results, especially the connection to maximally recoverable codes. The problem of zero-error computation of symmetric functions in a network setting is studied in [22], [23]. In [22], the authors characterize the rate at which symmetric functions (eg: mean, mode, max etc.) can be computed at a

sink node, while in [23], the authors provide optimal communication strategies for computing symmetric Boolean functions. The works of [24], [25], [26] study zero-error function computation in networks under the framework of network coding. In [24], the authors study the significance of the min-cut of the network when a destination node is interested in computing a general function of a set of independent source nodes. The works of [25], [26] consider sum networks, where a set of destination nodes in the network is interested in computing the sum of the observations corresponding to a set of source nodes.

The organization of the rest of the paper is as follows. Sections II contains definitions and facts relating to maximally recoverable codes. Necessary and sufficient conditions for achieving optimal communication cost in the point-to-point case appears in Section III. In Section IV, we present a low field size construction for the setting considered in Example 1. The broadcast setting and the associated problem of finding *sandwiched* MRSCs are discussed in Section V. Finally, our conclusions appear in Section VI.

Notation : Given a matrix $A \in \mathbb{F}_q^{m\times n}, m \leq n$ having rank m, we write $\mathcal{C}_A$ to denote the $[n, m]$ linear block code over $\mathbb{F}_q$ generated by the rows of A. The rank of matrix A will be denoted as $\rho(A)$. We write $\mathcal{C}_A^{\perp}$ to denote the dual code of $\mathcal{C}_A$. For any set $S \subseteq [n] = \{1, 2, \ldots, n\}$ and for any n-length linear code $\mathcal{C}$, we use $\mathcal{C}|_S$ to denote the restriction of the code $\mathcal{C}$ to the set of coordinates indexed by the set S. Also, we write $\mathcal{C}^S$ to denote the subcode of $\mathcal{C}$ obtained by *shortening* $\mathcal{C}$ to the set of coordinates indexed by the set S. In other words, $\mathcal{C}^S$ denotes the set of all those codewords in $\mathcal{C}$ whose *support* is confined to $\mathcal{S}$. For any codeword $\mathbf{c} \in \mathcal{C}$ such that $\mathbf{c} = [c_1\ c_2\ \cdots\ c_n]$, the support of $\mathbf{c}$ is defined as $\text{supp}(\mathbf{c}) = \{i \in [n], c_i \neq 0\}$. We recall the well known fact that $(\mathcal{C}|_S)^{\perp} = \left(\mathcal{C}^{\perp}\right)^S$.

II. Background on Maximally Recoverable Codes

In this section, we review relevant known facts regarding MRSCs, including equivalent definitions, existence of MRSCs and the notion of MRCs with locality.

Definition 2 (ℓ-cores [6]): Consider an $[n, k]$ code $\mathcal{C}$, and let $\mathcal{C}^{\perp}$ denote the dual of $\mathcal{C}$. Any set $S \subseteq [n], |S| = \ell$ is termed as an ℓ-core of $\mathcal{C}^{\perp}$ if $\text{supp}\,(\mathbf{c}) \nsubseteq S, \forall \mathbf{c} \in \mathcal{C}^{\perp}$.

We next state certain equivalent definitions of MRSCs. These are mostly known from various existing works in the context of MRCs with locality, and are straightforward to verify. A proof is however included in Appendix A for the sake of completeness.

Lemma 2.1: Consider an $[n, t]$ code $\mathcal{C}_0$, and let $\mathcal{C}$ denote an $[n, k]$ subcode of $\mathcal{C}$. Also let G_0 and G denote generator matrices for $\mathcal{C}_0$ and $\mathcal{C}$, respectively. Then, the following statements are equivalent:

1. $\mathcal{C}$ is a maximally recoverable subcode of $\mathcal{C}_0$, as defined in Definition 1, i.e., for any set $S \subseteq [n], |S| = k$ such that $\rho\,(G_0|_S) = k$, we have $\rho\,(G|_S) = k$.
2. Any set $S \subseteq [n], |S| = k$ which is a k-core of $\mathcal{C}_0^{\perp}$ is also a k-core of $\mathcal{C}^{\perp}$.
3. For any set $S \subseteq [n], |S| = k$ which is a k-core of $\mathcal{C}_0^{\perp}$, we have $\rho(H|_{[n]\setminus S}) = n - k$. Here H denotes a parity check matrix for the code $\mathcal{C}$.
4. For any set $S \subseteq [n], |S| \leq k$, $\rho\,(G_0|_S) = \rho\,(G|_S)$. Since the dual of a punctured code is a shortened code of the dual code, this is equivalent to saying that $\mathcal{C}$ is an MRSC of $\mathcal{C}_0$ if and only if for any $S \subseteq [n], |S| \leq k$, $\left(\mathcal{C}^{\perp}\right)^S = \left(\mathcal{C}_0^{\perp}\right)^S$. We note that this is further equivalent to saying that any k-sparse vector $\mathbf{c}$ that is a codeword of $\mathcal{C}^{\perp}$, is also a codeword of $\mathcal{C}_0^{\perp}$.

Proof: See Appendix A. ∎

The following lemma, which is restatement of Lemma 14 in [6], guarantees the existence of MRSCs under sufficiently large field size.

Lemma 2.2 ([6]): Given any $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, and any k such that $k < t$, there exists an $[n, k]$ maximally recoverable subcode $\mathcal{C}$ of $\mathcal{C}_0$, whenever $q > kn^k$.

We next present an explicit construction of MRSCs based on matrix representations corresponding to *linearized polynomials*. Linearized polynomial based constructions for MRCs with locality, from a parity-check matrix point

of view, appear in [5]. In the same paper [5], the authors show, again using linearized polynomials and parity-check matrix ideas, how to obtain an explicit construction of an MRSC $\mathcal{C}$ of $\mathcal{C}_0$, when G_0 is a binary matrix. The construction described below[3] uses a generator matrix view point, and the technique is similar to the usage of linearized polynomials appearing in works of [27], [28], [29]. These works used linearized polynomials for constructing locally repairable codes. As we will see, the construction of the *sandwiched* MRSCs that we present in the context of the broadcast setting (see Fig. 2) is an adaptation of the following construction.

Construction 2.3: Consider the code $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, having a generator matrix G_0. Let $\mathbb{F}_Q$ denote an extension field of $\mathbb{F}_q$, where $Q = q^t$. Also, let $\{\alpha_i \in \mathbb{F}_Q, 1 \leq i \leq t\}$ denote a basis of $\mathbb{F}_Q$ over $\mathbb{F}_q$. Define the elements $\{\beta_i \in \mathbb{F}_Q, 1 \leq i \leq n\}$ as follows:

$$[\beta_1 \ \beta_2 \ \cdots \ \beta_n] \ = \ [\alpha_1 \ \alpha_2 \ \cdots \ \alpha_t]G_0. \tag{12}$$

Next, consider the $[n, k]$ code $\mathcal{C}$ over $\mathbb{F}_Q$ having a generator matrix G given by

$$G \ = \ \begin{bmatrix} \beta_1 & \beta_2 & \cdots & \beta_n \\ \beta_1^q & \beta_2^q & \cdots & \beta_n^q \\ & & \vdots & \\ \beta_1^{q^{k-1}} & \beta_2^{q^{k-1}} & \cdots & \beta_n^{q^{k-1}} \end{bmatrix}. \tag{13}$$

The code $\mathcal{C}$ is the candidate MRSC of $\mathcal{C}_0$, where we consider $\mathcal{C}_0$ itself as a code over $\mathbb{F}_Q$. This is formally stated in the following theorem.

Theorem 2.4: Consider an $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, having a generator matrix G_0, where $G_0 \in \mathbb{F}_q^{t \times n}$. Let $\mathcal{C}_0^{(Q)}$ denote the $[n, t]$ code over $\mathbb{F}_Q$ that is also generated by G_0, where $Q = q^t$. Then the $[n, k]$ code $\mathcal{C}$ over $\mathbb{F}_Q$ obtained in Construction 2.3 is a maximally recoverable subcode of $\mathcal{C}_0^{(Q)}$.

Proof: See Appendix B. ∎

A. Maximally Recoverable Codes with Locality (Partial MDS Codes)

Below we give the definition of MRCs with locality, using the notion of MRSCs presented in Definition 1.

Definition 3 (MRCs with Locality [5], [6]): Assume that the parity matrix H_0 of $\mathcal{C}_0$ has the following form:

$$H_0 \ = \ \begin{bmatrix} H_{L,1} & & & \\ & H_{L,2} & & \\ & & \ddots & \\ & & & H_{L,\ell} \end{bmatrix}, \tag{14}$$

where each $H_{L,i}, 1 \leq i \leq \ell$ generates an $[r + \delta - 1, \delta - 1]$ MDS code, for some fixed r, δ. The code $\mathcal{C}_0^{\perp}$ has parameters $[n = \ell(r + \delta - 1), n - t = \ell(\delta - 1)]$. Then an $[n, k]$ MRSC $\mathcal{C}$ of $\mathcal{C}_0$ will be referred to as an $[n, k]$ (r, δ) maximally recoverable code with locality.

We would like to point out that the precise form of $H_{L,i}, 1 \leq i \leq \ell$ is not imposed by the above definition. In other words, while designing an $[n, k]$ (r, δ) MRC with locality, we are free to choose any $H_{L,i}, 1 \leq i \leq \ell$ that generates an $[r + \delta - 1, r, \delta]$ MDS code. MRCs with locality are also known in literature as partial MDS codes [4]. An $[n, k]$ (r, δ) MRC with locality corresponds to an $[m', n'](r', s')$ partial MDS code [4], where

$$m' \ = \ \frac{n}{r + \delta - 1} \tag{15}$$

$$n' \ = \ r + \delta - 1 \tag{16}$$

$$r' \ = \ \delta - 1 \tag{17}$$

$$s' \ = \ \frac{n}{r + \delta - 1} r - k. \tag{18}$$

[3]We do not claim novelty for Construction 2.3, since the idea used here follows directly from works like [27], [28], [29]

[4]Notation used for partial MDS codes correspond to the one used in [4].

The idea here is that we see the code as a two dimensional array code having m' rows and n' columns. Each row of the array is an $[n' = r + \delta - 1, n' - r' = r]$ MDS code having $r' = \delta - 1$ local parities. The parameter $s' = \dim(\mathcal{C}_0) - \dim(\mathcal{C})$ refers to the number of global parity symbols. An $[m', n'](r', s')$ partial MDS code can tolerate any combination of r' erasures in each row, and an additional s' erasures among the remaining elements in the array. In this write-up we will follow the terminology of MRCs with locality.

III. Communication Cost for the Point-to-Point Case

In this section, we obtain necessary and sufficient conditions for solving the point-to-point problem setting in Fig. 1. An encoder-decoder pair $(H, \mathcal{D})$ will be referred to as a *valid scheme* for the problem in Fig. 1, if the decoder's estimate of $A(X^n + E^n)$ is correct for all $X^n, E^n \in \mathbb{F}_q^n$, such that Hamming wt.$(E^n) \leq \epsilon$. Recall our assumption that we deal with a zero probability of error, worst-case scenario model. Also, recall that the communication cost associated with the encoder H is given by $\ell = \rho(H)$, where the parameters n, q, ϵ and A of the system are assumed to be fixed a priori. Without loss of generality, we assume that the $m \times n$ matrix A has rank m.

A. Necessary Conditions

Theorem 3.1: Consider any valid scheme $(H, \mathcal{D})$ for the communication problem described in Fig. 1. Let $\mathcal{C}_H$ and $\mathcal{C}_A$ denote the linear codes generated by the rows of the matrices H and A respectively. Also, consider the code $\mathcal{C} = \mathcal{C}_H \cap \mathcal{C}_A$, and let B denote a generator matrix for $\mathcal{C}$. Then, the following conditions must necessarily be satisfied:

1. If Y^n is any 2ϵ-sparse vector such that $AY^n \neq \mathbf{0}$, then it must also be true that $BY^n \neq \mathbf{0}$.
2. $\dim(\mathcal{C}) \geq \min(m, 2\epsilon)$.

Combining the two parts it then follows that if $\dim(\mathcal{C}) = \min(m, 2\epsilon)$, then $\mathcal{C}$ must be a maximally recoverable subcode of $\mathcal{C}_A$.

Proof: 1. We will prove the first part by contradiction. Let us suppose that there exists a non-zero 2ϵ-sparse vector Y^n such that $Y^n \in \mathcal{C}^\perp$, but $Y^n \notin \mathcal{C}_A^\perp$. In this case, we will show that there exists two distinct pairs (X_1^n, E_1^n) and (X_2^n, E_2^n), with E_1^n and E_2^n being ϵ-sparse, such that $H(X_1^n + E_1^n) = H(X_2^n + E_2^n)$ and $AX_1^n = AX_2^n$, but $AE_1^n \neq AE_2^n$. Thus, no decoder $\mathcal{D}$ can resolve between the two pairs successfully, and we would have contradicted the validity of our scheme. Towards constructing the desired (X^n, E^n) pairs, note that Y^n can be expressed as $Y^n = E_1^n - E_2^n$ for some two distinct ϵ-sparse vectors E_1^n and E_2^n. Also, since $\mathcal{C}^\perp = \mathcal{C}_H^\perp + \mathcal{C}_A^\perp$, there exists $U^n \in \mathcal{C}_A^\perp$ and $V^n \in \mathcal{C}_H^\perp$ such that $E_1^n - E_2^n = U^n + V^n$. We now choose $X_1^n = -U^n$ and $X_2^n = \mathbf{0}$. In this case, we see that $H(X_1^n + E_1^n) = H(X_2^n + E_2^n) = HE_2^n$ and $AX_1^n = AX_2^n = \mathbf{0}$, but $AE_1^n \neq AE_2^n$, which contradicts the validity of our scheme. This completes the proof of the first part.

2. Assume that $\dim(\mathcal{C}) < \min(m, 2\epsilon)$ which implies that $\dim\left(\mathcal{C}^\perp\right) > n - \min(m, 2\epsilon)$. In this case, it is straightforward to see that there exists a basis $\mathcal{L}$ of $\mathcal{C}^\perp$ consisting entirely of vectors that are 2ϵ-sparse. Also, note that $\dim(\mathcal{C}) < \min(m, 2\epsilon) \leq \dim(\mathcal{C}_A)$, which implies that $\dim\left(\mathcal{C}^\perp\right) > \dim\left(\mathcal{C}_A^\perp\right)$. Thus at least one of the elements of the set $\mathcal{L}$ is not contained in $\mathcal{C}_A^\perp$. In other words, there exists a non-zero 2ϵ-sparse vector Y^n such that $Y^n \in \mathcal{C}^\perp$, but $Y^n \notin \mathcal{C}_A^\perp$. The rest of the proof (to arrive at contradiction) follows from the proof of Part 1. of the theorem. We thus conclude that $\dim(C) \geq \min(m, 2\epsilon)$.

Finally, assume that $m > 2\epsilon$, and note that Part 1 implies that if $S \subseteq [n], |S| = 2\epsilon$ is such that $\rho(A|_S) = 2\epsilon$, then $\rho(B|_S) = 2\epsilon$. It then follows from Definition 1 that $\mathcal{C}$ must be a MRSC of $\mathcal{C}_A$, if $\dim(\mathcal{C}) = 2\epsilon$. ■

Corollary 3.2: a) The communication cost ℓ associated with any valid scheme for the problem given in Fig. 1 is lower bounded by $\ell \geq \min(m, 2\epsilon)$.

b) If $m \leq 2\epsilon$ and $\ell = m$ (optimal communication cost), then the encoding matrix H is necessarily given by $H = A$.

c) If $m > 2\epsilon$ and $\ell = 2\epsilon$ (once again optimal communication cost), then it must necessarily be true that the code $\mathcal{C}_H$ is a MRSC of $\mathcal{C}_A$.

Remark 1: The bound $\ell \geq \min(m, 2\epsilon)$ in Part a) of the above corollary was already proved in [9] for an average case setting, under the assumption that the vectors X^n and $X^n + E^n$ are uniformly picked at random from the set of all vectors satisfying Hamming wt.$(E^n) \leq \epsilon$. This result in [9] can be directly used to argue the correctness of Part a) of the above corollary, for the worse case setting that we consider here. However, the connection to maximally recoverable codes, especially the necessity of MRSCs for achieving optimality (Part c) of the above

corollary) is a novelty of this paper, and as we we will see next forms the basis of the our achievability scheme for an arbitray matrix A chosen for a sufficiently large field size.

B. Optimal Achievable Scheme based on Maximally Recoverable Subcodes, when $m > 2\epsilon$

We now present a valid scheme for the case $m > 2\epsilon$, having optimal communication cost $\ell = 2\epsilon$. From Corollary 3.2 , we know that the $[n, 2\epsilon]$ code $\mathcal{C}_H$ must necessarily be an MRSC of $\mathcal{C}_A$. We also know from Lemma 2.2 that such an MRSC always exists whenever the field size $q > 2\epsilon n^{2\epsilon}$. Since $\mathcal{C}_H$ is a subcode of $\mathcal{C}_A$, there exists a $2\epsilon \times m$ matrix S such that $H = SA$. Given the matrices H and S, we now refer to Fig. 3 for a schematic of the decoder to be used. The various steps performed by the decoder to estimate $A(X^n + E^n)$ are as follows:

(1) Given the encoder output $H(X^n + E^n)$ and the side information AX^n, obtain HE^n as

$$HE^n \;\; = \;\; H(X^n + E^n) - S\left(AX^n\right). \tag{19}$$

(2) Determine any ϵ-sparse vector $\widehat{E}^n$ such that $HE^n = H\widehat{E}^n$. To make the decoder deterministic, assume that $\widehat{E}^n$ is the least Hamming weight vector such that $HE^n = H\widehat{E}^n$. Since the vector $\left(E^n - \widehat{E}^n\right)$ is 2ϵ-sparse, and since $\mathcal{C}_H$ is a 2ϵ-dimensional MRSC of $\mathcal{C}_A$, we know from Part 4. of Lemma 2.1 that $HE^n = H\widehat{E}^n \implies AE^n = A\widehat{E}^n$.

(3) Compute the desired estimate as $AX^n + A\widehat{E}^n = A(X^n + E^n)$.

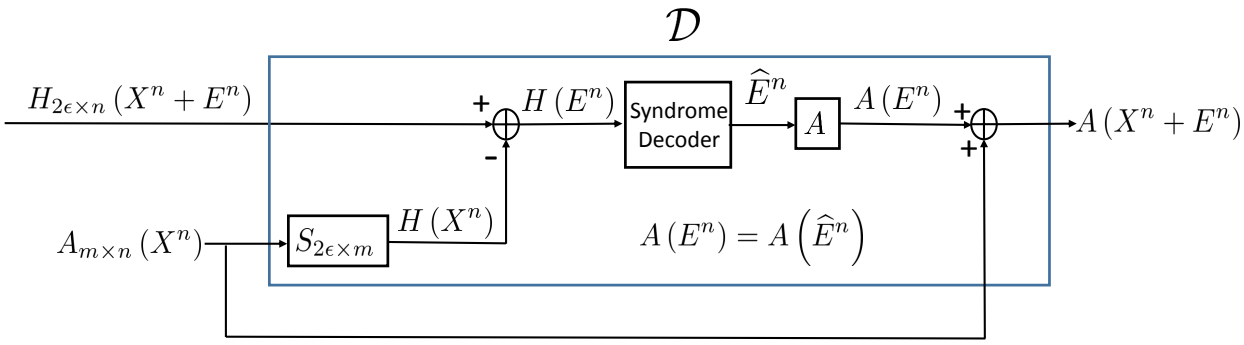

Fig. 3. The decoder used in the achievability proof for the case $m > 2\epsilon$.

Remark 2: We note that it is possible to construct a valid scheme, not necessarily optimal, using any matrix H which satisfies the conditions of Theorem 3.1. This also implies that if $(H, \mathcal{D})$ is any valid scheme for the problem, and if $\mathcal{C}_A \cap \mathcal{C}_H \subsetneq \mathcal{C}_A$, then a second valid scheme for the problem having lower communication cost can be constructed using the matrix B for encoding, where B is a generator matrix for the code $\mathcal{C}_A \cap \mathcal{C}_H$. This follows because the matrix B will also satisfy the conditions of the Theorem 3.1.

IV. Updating Linearly Encoded Striped-Data-Files

In this section, we present a low-field MRSC construction for the setting of Example 1, where we allow striping and encoding by an arbitrary $[N, K]$ linear code $\mathcal{C}_{st}$ over $\mathbb{F}_q$. Recall that in Example 1, we apply the point-to-point model in Fig. 1 in order to update the coded data in any one of the N storage nodes in the system. The K-length coding vector for one of the storage nodes is given by $\mathbf{a} = [a_1 \; a_2 \; \cdots \; a_K], a_i \in \mathbb{F}_q$. The vector X^n denotes the uncoded data file, which is divided into m stripes each of length K symbols. The first stripe consists of the first K symbols of X^n, the second stripe consists of symbols $X_{K+1}, \ldots, X_{2k}$, and so on. We assume that the length $n = mK$. The overall coding matrix A, which corresponds to the desired function at the destination, is given by

$$A \;\; = \;\; \begin{bmatrix} \mathbf{a} & & & \\ & \mathbf{a} & & \\ & & \ddots & \\ & & & \mathbf{a} \end{bmatrix}. \tag{20}$$

If we design the optimal encoder H based on existence of MRSCs from Lemma 2.2, then the field size q must be at least $2\epsilon n^{2\epsilon}$.

We now present an alternate low field size explicit construction of MRSCs of $\mathcal{C}_A$, when the matrix A takes on the special form given in (20). For the coding vector $\mathbf{a} = [a_1\ a_2\ \cdots\ a_K]$, if assume that $a_i \neq 0, \forall i \in [K]$, then note that an $[n, 2\epsilon]$ MRSC $\mathcal{C}_H$ of $\mathcal{C}_A$ is in fact an $[n, 2\epsilon](K, 1)$ MRC with locality (see Definition 3), i.e., a code with locality where all the local codes are scaled repetition codes. The problem of low field size constructions of MRCS with locality is in general an open problem. However, as we show here, for the special case when the local codes appear as repetition codes, low field size constructions are easily identified.

We divide the discussion into two parts. We will first show how MRSCs of any general code $\mathcal{C}_0$ can be obtained by first suitably *extending* $\mathcal{C}_0$, and then shortening the extended code. Given this result, we then show that the problem of constructing an MRSC of $\mathcal{C}_A$ where A is as given by (20), is as simple as finding an $[m, m-2\epsilon]$ MDS code over $\mathbb{F}_q$. If we employ Reed-Solomon codes, a field size $q > m$ is sufficient for the construction.

A. An Equivalent Definition of MRSCs via Code Extensions

Consider an $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$ having generator matrix G_0. Define a new $[n+\Delta, t]$ code $\mathcal{C}_0^{(e)}$ over $\mathbb{F}_q$ whose generator matrix $G_0^{(e)}$ is given by

$$G_0^{(e)} \quad = \quad [G_0 \mid Q], \tag{21}$$

for some $t \times \Delta$ matrix Q over $\mathbb{F}_q$. Code $\mathcal{C}_0^{(e)}$ will be referred to as extension of $\mathcal{C}_0$. In the following lemma, we show that MRSCs of $\mathcal{C}_0$ exist if and only if certain extensions of $\mathcal{C}_0$ exist.

Lemma 4.1: Consider an $[n+\Delta, t]$ extension $\mathcal{C}_0^{(e)}$ of the $[n, t]$ code $\mathcal{C}_0$ such that (i) $\Delta < t$, and (ii) for any $S \subseteq [n], |S| = t - \Delta$, we have

$$\rho\left(G_0^{(e)}|_{S\cup\{n+1,\ldots,n+\Delta\}}\right) \quad = \quad \rho\left(G_0|_S\right) + \Delta. \tag{22}$$

Then the shortened code $\left(\mathcal{C}_0^{(e)}\right)^{[n]}$ is an $[n, k = t-\Delta]$ MRSC of $\mathcal{C}_0$. Conversely, suppose that there exists an $[n, k]$ MRSC of the $[n, t]$ code $\mathcal{C}_0$, where $k = t - \Delta$ for some $\Delta < t$. Then, there exists an $[n+\Delta, t]$ extension $\mathcal{C}_0^{(e)}$ of the $[n, t]$ code $\mathcal{C}_0$ that satisfies (22) for any $S \subseteq [n], |S| = t - \Delta$.

Proof: We only prove the forward part here, a proof of converse appears in Appendix C. To prove the forward part, we assume the existence of extension $\mathcal{C}_0^{(e)}$ satisfying (22), and show that the shortened code $\mathcal{C} \triangleq \left(\mathcal{C}_0^{(e)}\right)^{[n]}$ is an $[n, k = t-\Delta]$ MRSC of $\mathcal{C}_0$. From Part 4, Lemma 2.1, we know it is sufficient to prove that

$$\left(\mathcal{C}^{\perp}\right)^S \quad = \quad \left(\mathcal{C}_0^{\perp}\right)^S \ \forall\ S \subseteq [n], |S| \leq k. \tag{23}$$

Let us suppose (23) is not true; i.e., there exist a set $S^* \subseteq [n], |S^*| \leq k$ such that $\left(\mathcal{C}_0^{\perp}\right)^{S^*} \subsetneq \left(\mathcal{C}^{\perp}\right)^{S^*}$. In other words, there exists a vector $\mathbf{c} \in \mathcal{C}^{\perp}$ such that $\mathbf{c} \notin \mathcal{C}_0^{\perp}$ and supp$(\mathbf{c}) \subseteq S^*$. Next, note that if H_0 denotes a parity check matrix for $\mathcal{C}_0$, then parity check matrices $H_0^{(e)}$ and H for the codes $\mathcal{C}_0^{(e)}$ and $\mathcal{C}$ can be respectively given by

$$H_0^{(e)} \ = \ \begin{bmatrix} H_0 & 0_{n-t\times\Delta} \\ H_e & I_\Delta \end{bmatrix} \quad \text{and} \quad H = \begin{bmatrix} H_0 \\ H_e \end{bmatrix}, \tag{24}$$

where H_e is some $\Delta \times n$ matrix over $\mathbb{F}_q$, and I_Δ denotes the $\Delta \times \Delta$ identity matrix. The fact that the parity-check matrix H for the code $\mathcal{C}$ is as given above, follows from the fact the dual of a shorten code is simply the punctured code of the dual code. In this case, corresponding to the vector $\mathbf{c}$, there exists a vector $\mathbf{c}^{(e)} \in \left(\mathcal{C}_0^{(e)}\right)^{\perp}$ such that

- supp $\left(\mathbf{c}^{(e)}\right) \subseteq S^* \cup \{n+1, \cdots, n+\Delta\}$,
- $\mathbf{c} = \left(\mathbf{c}^{(e)}\right)|_{S^*}$, and
- the set supp $\left(\mathbf{c}^{(e)}\right) \cap \{n+1, \cdots, n+\Delta\}$ is not empty. The last part follows from our assumption that $\mathbf{c} \notin \mathcal{C}_0^{\perp}$.

The presence of the vector $\mathbf{c}^{(e)} \in \left(\mathcal{C}_0^{(e)}\right)^{\perp}$ satisfying the above conditions means that

$$\rho\left(G_0^{(e)}|_{\mathcal{S}^*\cup\{n+1,\ldots,n+\Delta\}}\right) \quad < \quad \rho\left(G_0|_{\mathcal{S}^*}\right) + \Delta, \tag{25}$$

which is a contradiction to our assumption that the extension $\mathcal{C}_0^{(e)}$ satisfies (22). From this we conclude that (23) is indeed true, and this completes the proof of the forward part. ■

B. Low Field Size MRSCs for Updating Striped-Data-Files

Construction 4.2: Consider the matrix $A \in \mathbb{F}_q^{m\times n}$ as given in (20), and the associated $[n, m]$ linear code $\mathcal{C}_A$ over $\mathbb{F}_q$. Next consider the $[n + m - 2\epsilon, m]$ extended code $\mathcal{C}_A^{(e)}$ over $\mathbb{F}_q$, whose generator matrix $A^{(e)}$ is given by

$$A^{(e)} \;=\; [A \mid Q]\,, \tag{26}$$

where $Q \in \mathbb{F}_q^{m\times(m-2\epsilon)}$, and is such that Q^T generates an $[m, m - 2\epsilon]$ MDS code over $\mathbb{F}_q$. The candidate for the desired MRSC $\mathcal{C}_H$ is given by

$$\mathcal{C}_H \;=\; \left(\mathcal{C}_A^{(e)}\right)^{[n]}. \tag{27}$$

Note that $\mathcal{C}_H$ is an $[n, 2\epsilon]$ code. This completes the description of the construction.

Theorem 4.3: The $[n, 2\epsilon]$ code $\mathcal{C}_H$ obtained in Construction 4.2 is an MRSC of the $[n, m]$ code $\mathcal{C}_A$, where A is given in (20).

Proof: It is straightforward to check that the extended matrix $A^{(e)}$ in (26) satisfies the condition given in (22), i.e., for any $S \subseteq [n], |S| = 2\epsilon$, we have

$$\rho\left(A^{(e)}|_{S\cup\{n+1,\ldots,n+m-2\epsilon\}}\right) \;=\; \rho\left(A|_S\right) + (m - 2\epsilon). \tag{28}$$

The proof now follows from the forward part of Lemma 4.1. ■

Remark 3: The converse part of Lemma 4.1 can be used to show that given any $[n, 2\epsilon]$ MRSC $\mathcal{C}_H$ of $\mathcal{C}_A$, and if the vector $\mathbf{a}$ is such that $a_i \neq 0\ \forall i \in [K]$, then there exists an extended matrix $A^{(e)} = [A \mid Q]$ such that Q^T generates an $[m, m - 2\epsilon]$ MDS code over $\mathbb{F}_q$. Note that the assumption $a_i \neq 0$ was not imposed in Theorem 4.3.

V. Broadcasting to Receivers Interested in Different Functions

In the broadcast setting (see Fig. 2), we consider two receivers that are interested in computing two separate linear functions of the message vector. The functions correspond to the matrices A and B, respectively. Though the rest of the parameters are similar to those of the point-to-point case, we quickly describe them here again for sake of clarity. The vector $X^n \in \mathbb{F}_q^n$ denotes the initial source message. The two receivers hold AX^n and BX^n as side information, respectively. The matrices A and B have sizes $m_A \times n$ and $m_B \times n$, respectively, where $m_A, m_B \leq n$. Without loss of generality, we assume that $\rho(A) = m_A$ and $\rho(B) = m_B$. The updated source message is given by vector $X^n + E^n$, where the difference-vector E^n is ϵ-sparse. Encoding is carried out via the $\ell \times n$ matrix H. The goal is to recover the functions $A(X^n + E^n)$ and $B(X^n + E^n)$ at the respective receivers. The decoders used at the two receivers are denoted by $\mathcal{D}_1$ and $\mathcal{D}_2$, respectively. Once again, we assume a zero-probability-of-error, worst-case-scenario model. We assume knowledge of the functions A, B and the parameter ϵ, while designing the encoder H. The communication cost for the model is given by ℓ, assuming that the parameters n, q, ϵ, A and B are fixed. The triplet $(H, \mathcal{D}_1, \mathcal{D}_2)$ will be referred to as a *valid scheme* for the broadcast problem, if both the decoders' estimates are correct for all $X^n, E^n \in \mathbb{F}_q^n$ such that Hamming wt.$(E^n) \leq \epsilon$.

Our goal in this section is to identify necessary and sufficient conditions on valid schemes for the broadcast problem. Specifically, we are interested in characterizing the minimum communication cost (among valid schemes) that can be achieved for the setting. We divide the discussion into three parts. We first consider the special case when the two linear codes $\mathcal{C}_A$ and $\mathcal{C}_B$ intersect trivially, i.e., $\mathcal{C}_A \cap \mathcal{C}_B = \{\mathbf{0}\}$. The optimal communication cost for this case is straightforward to compute, given the observations from the point-to-point case. We then present necessary and sufficient conditions for the existence of sandwiched MRSCs. Recall that in the problem of sandwiched MRSCs, we begin with a code $\mathcal{C}_0$ and a subcode $\widehat{\mathcal{C}}$ of $\mathcal{C}_0$. The goal is to identify a MRSC $\mathcal{C}$ of $\mathcal{C}_0$ such that $\widehat{\mathcal{C}} \subset \mathcal{C}$. Finally, we will show how the concept of sandwiched MRSCs can be used to identify valid schemes having optimal communication cost for the case of arbitrary matrices A and B. Illustration of the results will be done by analyzing the settings in Examples 2 and 3.

A. Optimal Communication Cost for the Case $\mathcal{C}_A \cap \mathcal{C}_B = \{\mathbf{0}\}$

Let $(H, \mathcal{D}_1, \mathcal{D}_2)$ denote a valid scheme for the case $\mathcal{C}_A \cap \mathcal{C}_B = \{\mathbf{0}\}$. For decoders $\mathcal{D}_1$ and $\mathcal{D}_2$ to be successful, we know from Theorem 3.1 that $\dim(\mathcal{C}_A \cap \mathcal{C}_H) \geq \min(m_A, 2\epsilon)$ and $\dim(\mathcal{C}_B \cap \mathcal{C}_H) \geq \min(m_B, 2\epsilon)$, respectively. From this if follows that for the case when $\dim(\mathcal{C}_A + \mathcal{C}_B) = \dim(\mathcal{C}_A) + \dim(\mathcal{C}_B)$, the communication cost $\ell = \dim(\mathcal{C}_H)$ is lower bounded by

$$\ell \;\; \geq \;\; \min(m_A, 2\epsilon) + \min(m_B, 2\epsilon). \tag{29}$$

Given the achievability result from Section III-B, we see that the above bound is trivially achieved by encoding separately for the two receivers, where each of the two encodings is optimal for the respective receiver. We formally state the above observations in the following theorem:

Theorem 5.1: The optimal communication cost for the broadcast setting in Fig. 2 is given by

$$\ell \;\; = \;\; \min(m_A, 2\epsilon) + \min(m_B, 2\epsilon), \tag{30}$$

whenever the codes $\mathcal{C}_A$ and $\mathcal{C}_B$ intersect trivially. Achievability is guaranteed under the assumption that the field size $q > 2\epsilon n^{2\epsilon}$.

The following example illustrates the case under consideration.

Example 2 Revisited: Consider the setting in Example 2, where striping and encoding of a data file was done by an $[N, K]$ linear code over $\mathbb{F}_q$. Assume that the code is an $[N, K]$ MDS code with $K \geq 2$. We apply the broadcast model to simultaneously update the contents of two out of the N storage nodes. Also, recall our assumption that the K-length coding vectors associated with the two storage nodes are given by $\mathbf{a} = [a_1 \;\; a_2 \;\; \ldots a_K], \mathbf{b} = [b_1 \;\; b_2 \;\; \ldots b_K], a_i, b_i \in \mathbb{F}_q, i \in [K]$, and the overall coding matrices A and B corresponding to the m stripes are then given by

$$A \;\; = \;\; \begin{bmatrix} \mathbf{a} & & & \\ & \mathbf{a} & & \\ & & \ddots & \\ & & & \mathbf{a} \end{bmatrix} \quad \text{and} \quad B \;\; = \;\; \begin{bmatrix} \mathbf{b} & & & \\ & \mathbf{b} & & \\ & & \ddots & \\ & & & \mathbf{b} \end{bmatrix}, \tag{31}$$

where both A and B have size $m \times mK$. Under the assumption that the $[N, K]$ code is MDS with $K \geq 2$, it is straightforward to see that the codes $\mathcal{C}_A$ and $\mathcal{C}_B$ corresponding to the matrices A and B intersect trivially, i.e., $\mathcal{C}_A \cap \mathcal{C}_B = \{\mathbf{0}\}$. In this case, from Theorem 5.1, we know that broadcasting does not help; the source must transmit as though it is individually updating the two storage nodes.

B. Sandwiched Maximally Recoverable Subcodes

We now take a slight detour, and identify necessary and sufficient conditions for the existence of sandwiched MRSCs. We assume that we are given an $[n, t]$ code $\mathcal{C}_0$ and an $[n, s]$ subcode $\widehat{\mathcal{C}}$ of $\mathcal{C}_0$. The question that we are interested is whether we can find an $[n, k]$ MRSC $\mathcal{C}$ of $\mathcal{C}_0$ such that $\widehat{\mathcal{C}}$ is a subcode of $\mathcal{C}$. We assume that the parameters s, k, t satisfy the relation $s \leq k \leq t$. Also, let G, G_0 and $\widehat{G}$ denote generator matrices for the codes $\mathcal{C}$, $\mathcal{C}_0$ and $\widehat{\mathcal{C}}$, respectively. Recall from Definition 1 that for $\mathcal{C}$ to be an MRSC of $\mathcal{C}_0$, it must be true that

$$\rho(G_0|_S) = k \quad \Longrightarrow \quad \rho(G|_S) = k, \;\; S \subseteq [n], |S| = k. \tag{32}$$

Thus, a *necessary condition for the existence of the sandwiched MRSC* $\mathcal{C}$ can be given as follows:

$$\rho(G_0|_S) = k \quad \Longrightarrow \quad \rho\left(\widehat{G}|_S\right) = s, \;\; S \subseteq [n], |S| = k. \tag{33}$$

Note that the condition in (33) is equivalent to saying that $\rho\left(\widehat{G}|_S\right) = s$ for any $S \subseteq [n], |S| = k$ which is an k-core of $\mathcal{C}_0^{\perp}$. In the following lemma, we show the sufficiency of the condition in (33) for the existence of sandwiched MRSCs, under the assumption of a large enough finite field size q.

Lemma 5.2: Suppose that we are given an $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, and an $[n, s]$ subcode $\widehat{\mathcal{C}}$, whose generator matrices G_0 and $\widehat{G}$ satisfy (33), where k is such that $s < k < t$. Then, there exists an $[n, k]$ code $\mathcal{C}$ such that

- $\widehat{\mathcal{C}} \subseteq \mathcal{C}$ and
- $\mathcal{C}$ is a maximally recoverable subcode of $\mathcal{C}_0$,

whenever $q > \binom{n}{k}$.

Proof: See Appendix D. ■

We next present an alternate, explicit construction of sandwiched MRSCs appears using linearized polynomials.

Construction 5.3: Consider the code $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, having a generator matrix G_0. Also, consider the $[n, s]$ subcode $\widehat{\mathcal{C}}$ of $\mathcal{C}_0$, having a generator matrix $\widehat{G}$. Without loss of generality, assume that the generator matrix G_0 is given by

$$G_0 \;=\; \left[\begin{array}{c} \widehat{G} \\ B \end{array}\right], \tag{34}$$

for some $t-s \times n$ matrix B. Note that we have $\rho(B) = t-s$. Let $\mathbb{F}_Q$ denote an extension field of $\mathbb{F}_q$, where $Q = q^{t-s}$. Further, let $\{\alpha_i \in \mathbb{F}_Q, 1 \leq i \leq t-s\}$ denote a basis of $\mathbb{F}_Q$ over $\mathbb{F}_q$. Define the elements $\{\beta_i \in \mathbb{F}_Q, 1 \leq i \leq n\}$ as follows:

$$[\beta_1 \; \beta_2 \; \cdots \; \beta_n] \;=\; [\alpha_1 \; \alpha_2 \; \cdots \; \alpha_{t-s}]B. \tag{35}$$

Next, consider the code $\mathcal{C}$ over $\mathbb{F}_Q$ having a generator matrix G given by

$$G \;=\; \left[\begin{array}{cccc} \multicolumn{4}{c}{\widehat{G}} \\ \hline \\ \beta_1 & \beta_2 & \cdots & \beta_n \\ \beta_1^q & \beta_2^q & \cdots & \beta_n^q \\ & & \vdots & \\ \beta_1^{q^{k-s-1}} & \beta_2^{q^{k-s-1}} & \cdots & \beta_n^{q^{k-s-1}} \end{array}\right]. \tag{36}$$

The code $\mathcal{C}$ is the candidate for the $[n, k]$ sandwiched MRSC of $\mathcal{C}_0$, where we consider $\mathcal{C}_0$ and $\widehat{\mathcal{C}}$ themselves as codes over $\mathbb{F}_Q$. This is formally stated in the following theorem.

Theorem 5.4: Consider an $[n, t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, having a generator matrix $G_0 \in \mathbb{F}_q^{t \times n}$. Also, consider the $[n, s]$ subcode $\widehat{\mathcal{C}}$ of $\mathcal{C}_0$, having a generator matrix $\widehat{G}$ such that the following condition is satisfied:

$$\rho(G_0|_S) = k \quad \Longrightarrow \quad \rho\left(\widehat{G}|_S\right) = s, \; S \subseteq [n], |S| = k, \tag{37}$$

where k is such that $s \leq k \leq t$. Next, let $\mathcal{C}_0^{(Q)}$ and $\widehat{\mathcal{C}}^{(Q)}$ denote the $[n, t]$ and $[n, s]$ codes over $\mathbb{F}_Q$ that are respectively generated by G_0 and $\widehat{G}$, where $Q = q^{t-s}$. Then the $[n, k]$ code $\mathcal{C}$ over $\mathbb{F}_Q$ obtained in Construction 5.3 is a maximally recoverable subcode of $\mathcal{C}_0^{(Q)}$ such that $\widehat{\mathcal{C}}^{(Q)} \subseteq \mathcal{C}$.

Proof: See Appendix E. ■

C. Optimal Communication Cost for Arbitrary Matrices A and B under the Broadcast Setting

We now use the concept of sandwiched MRSCs and characterize the optimal communication cost for the broadcast problem when the codes $\mathcal{C}_A$ and $\mathcal{C}_B$ have a non-trivial intersection. Below, we first give a qualitative description of our approach, before presenting technical details. Consider a valid scheme $(H, \mathcal{D}_1, \mathcal{D}_2)$ that we design for the problem. Let H_A and H_B denote the rows in the row-space of H which help the decoders $\mathcal{D}_1$ and $\mathcal{D}_2$ recover $A(X^n + E^n)$ and $B(X^n + E^n)$, respectively. Also let $\mathcal{C}_{H_A}$ and $\mathcal{C}_{H_B}$ denote the linear codes generated by H_A and H_B respectively. Our approach to minimizing the communication cost is to pick the encoder H such that

- $\mathcal{C}_{H_A}$ is a 2ϵ-dimensional (assuming that $m_A > 2\epsilon$) MRSC of $\mathcal{C}_A$,
- $\mathcal{C}_{H_B}$ is a 2ϵ-dimensional (assuming that $m_B > 2\epsilon$) MRSC of $\mathcal{C}_B$, and
- also "maximize" the dimension of intersection between the subcodes $\mathcal{C}_{\mathcal{H}_\mathcal{A}}$ and $\mathcal{C}_{H_B}$. The extent to which $\dim(\mathcal{C}_{\mathcal{H}_\mathcal{A}} \cap \mathcal{C}_{\mathcal{H}_\mathcal{B}})$ can be maximized depends on the dimension of the intersection between $\mathcal{C}_\mathcal{A}$ and $\mathcal{C}_\mathcal{B}$.

For ease of presentation, we give the expression and proofs for optimal communication cost under the assumption[5] that both m_A and m_B are greater than 2ϵ. Under this assumption, consider the code $\widetilde{\mathcal{C}} = \mathcal{C}_A \cap \mathcal{C}_B$ having a generator matrix $\widetilde{H}$. Further, consider the quantities θ_A, θ_B and θ defined as follows:

$$\theta_A = \min_{\substack{S\subseteq[n],|S|=2\epsilon \\ S \text{ is a } 2\epsilon\text{-core of } \mathcal{C}_A^\perp}} \rho(\widetilde{H}|_S), \tag{38}$$

$$\theta_B = \min_{\substack{S\subseteq[n],|S|=2\epsilon \\ S \text{ is a } 2\epsilon\text{-core of } \mathcal{C}_B^\perp}} \rho(\widetilde{H}|_S), \text{ and} \tag{39}$$

$$\theta = \min(\theta_A, \theta_B). \tag{40}$$

Note that the parameters θ_A, θ_B and θ are entirely determined given the matrices A and B, and the sparsity parameter ϵ. The following theorem characterizes the optimal communication cost in terms of the parameter θ.

Theorem 5.5: The optimal communication cost for the broadcast setting shown in Fig. 2 is given by

$$\ell = 4\epsilon - \theta, \tag{41}$$

whenever m_A and m_B are both greater than 2ϵ, and where the parameter θ is as defined by (40). Achievability is guaranteed under the assumption that the field size $q > \max\left(\binom{n}{2\epsilon}, \binom{n}{\theta}\right)$.

Proof: Let us prove the converse first, i.e., we show that the communication cost of an encoder H associated with any valid scheme is lower bounded by $\ell = \rho(H) \geq 4\epsilon - \theta$. Towards this, consider the codes $\mathcal{C}_H$, $\mathcal{C}_A$, $\mathcal{C}_B$, and define the codes $\mathcal{C}_{H_A}$, $\mathcal{C}_{H_B}$ and $\widehat{\mathcal{C}}$ as follows:

$$\mathcal{C}_{H_A} = \mathcal{C}_H \cap \mathcal{C}_A, \tag{42}$$

$$\mathcal{C}_{H_B} = \mathcal{C}_H \cap \mathcal{C}_B, \text{ and} \tag{43}$$

$$\widehat{\mathcal{C}} = \mathcal{C}_{H_A} \cap \mathcal{C}_{H_B}. \tag{44}$$

Note that $\widehat{\mathcal{C}} \subseteq \widetilde{\mathcal{C}} = \mathcal{C}_A \cap \mathcal{C}_B$. Also assume that H_A, H_B and $\widehat{H}$ denote generator matrices for the codes $\mathcal{C}_{H_A}$, $\mathcal{C}_{H_B}$ and $\widehat{\mathcal{C}}$, respectively. From Theorem 3.1, we know that if Y^n is any 2ϵ-sparse vector such that $AY^n \neq \mathbf{0}$, then $H_A Y^n \neq \mathbf{0}$. From this it follows that

$$\rho(A|_S) = 2\epsilon \implies \rho(H_A|_S) = 2\epsilon, \ S \subseteq [n], |S| = 2\epsilon. \tag{45}$$

Next, consider the definition of θ_A in (38), and let $S^* \subseteq [n], |S^*| = 2\epsilon$ denote a 2ϵ-core of $\mathcal{C}_A^\perp$ such that

$$\rho(\widetilde{H}|_{S^*}) = \theta_A. \tag{46}$$

Now, consider the following chain of inequalities:

$$2\epsilon \overset{(a)}{=} \rho(A|_{S^*}) \tag{47}$$

$$\overset{(b)}{=} \rho(H_A|_{S^*}) \tag{48}$$

$$\overset{(c)}{\leq} \rho(\widehat{H}|_{S^*}) + \left(\rho(H_A) - \rho(\widehat{H})\right) \tag{49}$$

$$\overset{(d)}{\leq} \rho(\widetilde{H}|_{S^*}) + \left(\rho(H_A) - \rho(\widehat{H})\right) \tag{50}$$

$$\overset{(e)}{=} \theta_A + \left(\rho(H_A) - \rho(\widehat{H})\right). \tag{51}$$

Here, (a) follows because S^* is a 2ϵ-core of $\mathcal{C}_A^\perp$,
(b) follows from (45),
(c) follows by observing since $\widehat{\mathcal{C}}$ is a subcode of $\mathcal{C}_A$, we have $\rho(H_A) - \rho(\widehat{H}) \geq \rho(H_A|_S) - \rho(\widehat{H}|_S)$, for any set $S \subset [n]$. In particular, it is true that $\rho(H_A) - \rho(\widehat{H}) \geq \rho(H_A|_{S^*}) - \rho(\widehat{H}|_{S^*})$,

[5]This is the hardest of all the cases. The remaining cases when one or both the ranks m_A and m_B are less than 2ϵ can be similarly handled.

(d) follows since $\widehat{\mathcal{C}}$ is a subcode of $\widetilde{\mathcal{C}}$, and finally
(e) follows from (46).
In other words, we get that

$$\rho(\widehat{H}) \quad \leq \quad \theta_A + \rho(H_A) - 2\epsilon. \tag{52}$$

In a similar fashion, one can also show that

$$\rho(\widehat{H}) \quad \leq \quad \theta_B + \rho(H_B) - 2\epsilon. \tag{53}$$

The communication cost ℓ associated with the encoder H can now be lower bounded as follows:

$$\begin{aligned}
\ell &= \rho(H) && (54)\\
&\geq \rho(H_A) + \rho(H_B) - \rho(\widehat{H}) && (55)\\
&\overset{(a)}{\geq} \rho(H_A) + \rho(H_B) - \min(\theta_A + \rho(H_A) - 2\epsilon, \theta_B + \rho(H_B) - 2\epsilon) && (56)\\
&\geq \rho(H_A) + \rho(H_B) - \min(\theta_A, \theta_B) - \max(\rho(H_A), \rho(H_B)) + 2\epsilon && (57)\\
&\overset{(b)}{\geq} 4\epsilon - \min(\theta_A, \theta_B) = 4\epsilon - \theta. && (58)
\end{aligned}$$

Here (a) follows from (52) and (53), and (b) follows from our assumption that the ranks of both H_A and H_B are greater than or equal to 2ϵ. This completes the proof of the lower bound on the communication cost.

Proof of Achievability: Let us now show that it is indeed possible to construct a valid scheme having communication cost $\ell = 4\epsilon - \theta$, under the assumption of a sufficiently large field size. Towards this, consider the code $\widetilde{\mathcal{C}}$ and let $\widehat{\mathcal{C}}$ denote an θ-dimensional MRSC of $\widetilde{\mathcal{C}}$. We know from Lemma 5.2 that the code $\widehat{\mathcal{C}}$ always exists[6] whenever the field size $q > \binom{n}{\theta}$. Also, let $\widehat{H}$ denote a generator matrix for the code $\widehat{\mathcal{C}}$. Now, observe that if S is a 2ϵ-core of either $\mathcal{C}_A^{\perp}$ or $\mathcal{C}_B^{\perp}$, we know from the definition of θ in (40) that $\rho(\widetilde{H}|_S) \geq \theta$. Noting that $\theta \leq 2\epsilon$ and using the fact that $\widehat{\mathcal{C}}$ is an θ-dimensional MRSC of $\widetilde{\mathcal{C}}$, we get that $\rho(\widehat{H}|_S) = \theta$. In this case, we know from Lemma 5.2 that it is possible to identify a 2ϵ-dimensional MRSC $\mathcal{C}_{H_A}$ of $\mathcal{C}_A$ such that $\widehat{\mathcal{C}} \subseteq \mathcal{C}_{H_A}$, whenever the field size $q > \binom{n}{2\epsilon}$. Similarly, we can also identify a 2ϵ-dimensional MRSC $\mathcal{C}_{H_B}$ of $\mathcal{C}_B$ such that $\widehat{\mathcal{C}} \subseteq \mathcal{C}_{H_B}$, whenever the field size $q > \binom{n}{2\epsilon}$. The overall field size requirement, when we take into the account the minimum q needed for the existence of $\widehat{\mathcal{C}}$ is given by $q > \max\left(\binom{n}{2\epsilon}, \binom{n}{\theta}\right)$. The candidate code for the encoder is now given by $\mathcal{C}_H = \mathcal{C}_{H_A} + \mathcal{C}_{H_B}$. The communication cost of the encoder H is then given by $\ell = \text{rank}(H_A) + \text{rank}(H_B) - \dim(\mathcal{C}_{H_A} \cap \mathcal{C}_{H_B}) \leq \text{rank}(H_A) + \text{rank}(H_B) - \text{rank}(\widehat{H}) = 4\epsilon - \theta$. Also, we know from the achievability result of the point-to-point setting in Section III-B that decoders $\mathcal{D}_1$ and $\mathcal{D}_2$ in Fig. 2 can be constructed based on the matrices H_A and H_B respectively. This completes the proof of the achievability part of the theorem. ∎

Remark 4: In the above proof of achievability, we noted that $\widehat{\mathcal{C}} \subseteq \mathcal{C}_{H_A} \cap \mathcal{C}_{H_B}$. In fact, it is straightforward to see that $\widehat{\mathcal{C}} = \mathcal{C}_{H_A} \cap \mathcal{C}_{H_B}$; else we would contradict the minimality of either θ_A or θ_B or both.

The following example illustrates the case under consideration.

Example 3 revisited: We now revisit Example 3 where we considered striping and encoding of a data file by a $[N = 5, K = 3, D = 4](\alpha = 4, \beta = 1)$ MBR code over $\mathbb{F}_q$ that encodes 9 symbols into 20 symbols and stores across $N = 5$ nodes such that each node holds $\alpha = 4$ symbols. Recall that the contents of any $K = 3$ nodes are sufficient to reconstruct the 9 uncoded symbols, and that the contents of any one of the $N = 5$ node can be recovered by connecting to any set of $D = 4$ other nodes. Also, recall that the encoding is described by a

[6]Note that this follows from Lemma 2.2 as well; our use of Lemma 5.2 (with $\widehat{\mathcal{C}} = \{\mathbf{0}\}$ in Lemma 5.2) must be considered as a matter of choice.

product-matrx construction as follows:

$$C_{N\times\alpha} = \Psi_{N\times D} M_{D\times\alpha} \tag{59}$$

$$= \begin{bmatrix} \underline{\psi}_1 \\ \underline{\psi}_2 \\ \underline{\psi}_3 \\ \underline{\psi}_4 \\ \underline{\psi}_5 \end{bmatrix} \begin{bmatrix} m_1 & m_2 & m_3 & m_7 \\ m_2 & m_4 & m_5 & m8 \\ m_3 & m_5 & m_6 & m_9 \\ m_7 & m_8 & m_9 & 0 \end{bmatrix}, \tag{60}$$

where the encoding matrix Ψ can be chosen as a Vandermonde matrix given as follows:

$$\Psi = \begin{bmatrix} 1 & \gamma & \gamma^2 & \gamma^3 \\ 1 & \gamma^2 & \gamma^4 & \gamma^6 \\ 1 & \gamma^3 & \gamma^6 & \gamma^9 \\ 1 & \gamma^4 & \gamma^8 & \gamma^{12} \\ 1 & \gamma^5 & \gamma^{10} & \gamma^{15} \end{bmatrix}, \tag{61}$$

Here we pick γ as a primitive element in $\mathbb{F}_q$. The $N \times \alpha$ matrix C is the codeword matrix, with the i^{th} row representing the contents that get stored in the i^{th} node, $1 \leq i \leq 5$. It is known that any two nodes of an $[N = 5, K = 3, D = 4](\alpha = 4, \beta = 1)$ MBR code share exactly one non-trivial linear combination of the respective symbols. For example, in the case of the product-matrix MBR code considered here, the first two nodes share the linear combination

$$\underline{\psi}_1 M \underline{\psi}_2^t = \underline{\psi}_2 M \underline{\psi}_1^t. \tag{62}$$

The left-hand and the right-hand sides of (62) are linear combinations of the contents of the first and second nodes respectively, and the equality follows from the fact that the matrix M is symmetric.

Let us further recall the structure of the $4m \times 9m$ coding matrices A and B for the first two nodes, which are given by

$$A = \begin{bmatrix} A_1 & & & \\ & A_1 & & \\ & & \ddots & \\ & & & A_1 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} B_1 & & & \\ & B_1 & & \\ & & \ddots & \\ & & & B_1 \end{bmatrix}, \tag{63}$$

where

$$A_1 = \begin{bmatrix} 1 & \gamma & \gamma^2 & & & & \gamma^3 & & \\ & 1 & & \gamma & \gamma^2 & & & \gamma^3 & \\ & & 1 & & \gamma & \gamma^2 & & & \gamma^3 \\ & & & & & & 1 & \gamma & \gamma^2 \end{bmatrix} \quad \text{and} \quad B_1 = \begin{bmatrix} 1 & \gamma^2 & \gamma^4 & & & & \gamma^6 & & \\ & 1 & & \gamma^2 & \gamma^4 & & & \gamma^6 & \\ & & 1 & & \gamma^2 & \gamma^4 & & & \gamma^6 \\ & & & & & & 1 & \gamma^2 & \gamma^4 \end{bmatrix}. \tag{64}$$

Communication cost for updating two storage nodes via broadcasting : Let $\mathcal{C}_{A_1}$ and $\mathcal{C}_{B_1}$ denote the codes generated by the rows of A_1 and B_1. From (62), we know that $\dim(\mathcal{C}_{A_1} \cap \mathcal{C}_{B_1}) = 1$. It is straightforward to see that $\mathcal{C}_{A_1} \cap \mathcal{C}_{B_1}$ is generated by the vector

$$\mathbf{c} = [1,\ \gamma+\gamma^2,\ \gamma^2+\gamma^4,\ \gamma^3,\ \gamma^4+\gamma^5,\ \gamma^6,\ \gamma^3+\gamma^6,\ \gamma^5+\gamma^7,\ \gamma^7+\gamma^8]. \tag{65}$$

Note that if we satisfy the conditions

$$\begin{aligned} 1+\gamma &\neq 0 \\ 1+\gamma^2 &\neq 0 \\ 1+\gamma^3 &\neq 0 \end{aligned} \tag{66}$$

then, all the entires of the vector $\mathbf{c}$ are non-zeros. The conditions in (66) are trivially satisfied for the field sizes that are needed for the broadcast model (for e.g., even $q \geq 8$ is sufficient, since γ is primitive.). Next, consider the

codes $\mathcal{C}_A$ and $\mathcal{C}_B$ generated by the rows of A and B, and observe that a generator matrix $\widehat{G}$ for the intersection $\mathcal{C}_A \cap \mathcal{C}_B$ is given by

$$\widehat{G}_{m\times 9m} \;=\; \begin{bmatrix} \mathbf{c} & & & \\ & \mathbf{c} & & \\ & & \ddots & \\ & & & \mathbf{c} \end{bmatrix}. \tag{67}$$

We are now ready to calculate the quantities θ_A, θ_B and θ in (38)-(40), and further apply Theorem 5.5 to calculate the communication cost for broadcasting. We assume that $m > 2\epsilon$. It is straightforward to see that the quantities θ_A, θ_B and θ are given as follows:

$$\theta_A \;=\; \min_{\substack{S\subseteq[n],|S|=2\epsilon \\ S \text{ is a } 2\epsilon\text{-core of } \mathcal{C}_A^{\perp}}} \rho(\widetilde{H}|_S) \;=\; \left\lceil \frac{2\epsilon}{4} \right\rceil, \tag{68}$$

$$\theta_B \;=\; \min_{\substack{S\subseteq[n],|S|=2\epsilon \\ S \text{ is a } 2\epsilon\text{-core of } \mathcal{C}_B^{\perp}}} \rho(\widetilde{H}|_S) \;=\; \left\lceil \frac{2\epsilon}{4} \right\rceil, \tag{69}$$

$$\theta \;=\; \min(\theta_A, \theta_B) \;=\; \left\lceil \frac{2\epsilon}{4} \right\rceil. \tag{70}$$

The communication cost for broadcasting is then given by $\ell = 4\epsilon - \theta \approx 3.5\epsilon$. Note that if we individually update the two nodes, the overall communication will be 4ϵ. This completes our example.

VI. Conclusions

To conclude, in this paper, we considered a zero-error function update problem motivated by practical distributed storage systems, in which coded elements need to be updated by users whose data is out of sync with the data stored in the system. We studied the problem in its generality, by considering any possible linear update function. Optimal schemes for minimizing communication cost were obtained for point-to-point as well as broadcast settings. Examples based on well-known distributed storage codes were used to illustrate the potential applicability of the theory that we developed for both the point-to-point and broadcast cases. The most interesting contribution of this paper is perhaps the connection of the optimal solutions for the function update problem to the notion of maximally recoverable codes, which were conventionally studied as codes suitable for data storage applications. The current work provides an alternate view on this important class of codes, for which low field size constructions are still perused actively by many researchers. We introduced the class of sandwiched maximally recoverable codes, and showed its importance in constructing optimal schemes for the broadcast setting.

The work opens up several interesting practically motivated questions: 1) Given the fact the low field-size constructions of maximally recoverable codes are still largely unknown, can we obtain schemes for the function update problem which have sub-optimal communication costs, but have improved field-sizes ? In this regard, it will be interesting to explore connections with approximate variations of MRCs like sector-disk codes [4], [15], [16] or partial maximally recoverable codes [18]. 2) Can we enhance the system model to one which does not have exact knowledge of the the sparsity parameter ϵ? Recall that ϵ indicates the amount of update the source message undergoes. For instance, the source could encode assuming an approximate value of ϵ, and also send out a low-cost hash function which could be used to detect failure of decoding at the receiver. Assuming a feed-back link back to the source, we can perform further rounds of communication that incrementally update the destination function until decoding succeeds. 3) Can we use the results in this work to build secure incremental digital signature schemes, perhaps based on the McEliece crypto system [30]? We note that in Fig. 1, if we assume that the source message vector and update vector are arbitrary, the encoder's output does not seem to reveal much about either of these vectors.

References

[1] A.G. Dimakis, P.B. Godfrey, Y. Wu, M.J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *Information Theory, IEEE Transactions on*, vol. 56, no. 9, pp. 4539–4551, Sept 2010.

[2] K. V. Rashmi, N. B Shah, and P. V. Kumar, "Optimal exact-regenerating codes for distributed storage at the msr and mbr points via a product-matrix construction," *Information Theory, IEEE Transactions on*, vol. 57, no. 8, pp. 5227–5239, 2011.

[3] M. Chen, C. Huang, and J. Li, "On the maximally recoverable property for multi-protection group codes," in *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, June 2007, pp. 486–490.

[4] M. Blaum, J. L. Hafner, and S. Hetzler, "Partial-mds codes and their application to raid type of architectures," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4510–4519, July 2013.

[5] P. Gopalan, C. Huang, B. Jenkins, and S. Yekhanin, "Explicit maximally recoverable codes with locality," *IEEE Trans. Inf. Theory*, vol. 60, no. 9, pp. 5245–5256, Sept 2014.

[6] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the Locality of Codeword Symbols," *IEEE Trans. Inf. Theory*, vol. 58, no. 11, pp. 6925–6934, Nov. 2012.

[7] D. S. Papailiopoulos and A. G. Dimakis, "Locally repairable codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, MA, Jul. 2012, pp. 2771–2775.

[8] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, "Erasure coding in windows azure storage," in *Proc. USENIX Annual Technical Conference (ATC)*, 2012, pp. 15–26.

[9] P. Nakkiran, N. B. Shah, and K. V. Rashmi, "Fundamental limits on communication for oblivious updates in storage networks," in *2014 IEEE Global Communications Conference*, Dec 2014, pp. 2363–2368.

[10] P. Nakkiran, N. B. Shah, K. V. Rashmi, A. Sahai, and K. Ramchandran, "Optimal oblivious updates in distributed storage networks." [Online]. Available: http://people.eecs.berkeley.edu/~nihar/publications/ObUp.pdf

[11] S. E. Rouayheb, S. Goparaju, H. M. Kiah, and O. Milenkovic, "Synchronizing edits in distributed storage networks," in *Information Theory Proceedings (ISIT), 2015 IEEE International Symposium on*, June 2015.

[12] N. Ma, K. Ramchandran, and D. Tse, "Efficient file synchronization: A distributed source coding approach," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, July 2011, pp. 583–587.

[13] Q. Wang, V. R. Cadambe, S. Jaggi, M. Schwartz, and M. Médard, "File updates under random/arbitrary insertions and deletions," *CoRR*, vol. abs/1502.07830, 2015.

[14] A. Orlitsky and K. Viswanathan, "One-way communication and error-correcting codes," *Information Theory, IEEE Transactions on*, vol. 49, no. 7, pp. 1781–1788, July 2003.

[15] M. Blaum, J. S. Plank, M. Schwartz, and E. Yaakobi, "Partial mds (pmds) and sector-disk (sd) codes that tolerate the erasure of two random sectors," in *Information Theory (ISIT), 2014 IEEE International Symposium on*, June 2014, pp. 1792–1796.

[16] J. Chen, K. W. Shum, Q. Yu, and C. W. Sung, "Sector-disk codes and partial mds codes up to three global parities," in *Information Theory Proceedings (ISIT), 2015 IEEE International Symposium on*, June 2015.

[17] Mingqiang Li and Patrick P. C. Lee, "Stair codes: A general family of erasure codes for tolerating device and sector failures in practical storage systems," in *Proceedings of the 12th USENIX Conference on File and Storage Technologies (FAST 14)*. Santa Clara, CA: USENIX, 2014, pp. 147–162.

[18] S. B. Balaji and P. V. Kumar, "On partial maximally-recoverable and maximally-recoverable codes," in *Information Theory Proceedings (ISIT), 2015 IEEE International Symposium on*, June 2015.

[19] V. Lalitha and S. V. Lokam, "Weight enumerators and higher support weights of maximally recoverable codes," *CoRR*, vol. abs/1507.01095, 2015. [Online]. Available: http://arxiv.org/abs/1507.01095

[20] N. Prakash, V. Lalitha, and P. V. Kumar, "Codes with locality for two erasures," in *Information Theory (ISIT), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 1962–1966.

[21] H. Kowshik and P. R. Kumar, "Zero-error function computation in sensor networks," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, Dec 2009, pp. 3787–3792.

[22] A. Giridhar and P. R. Kumar, "Computing and communicating functions over sensor networks," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 4, pp. 755–764, April 2005.

[23] H. Kowshik and P. R. Kumar, "Optimal computation of symmetric boolean functions in collocated networks," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 4, pp. 639–654, April 2013.

[24] R. Appuswamy, M. Franceschetti, N. Karamchandani, and K. Zeger, "Network coding for computing: Cut-set bounds," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 1015 –1030, Feb. 2011.

[25] B. K. Rai and B. K. Dey, "On network coding for sum-networks," *IEEE Trans. Inf. Theory*, vol. 58, no. 1, pp. 50 –63, Jan. 2012.

[26] A. Ramamoorthy and M. Langberg, "Communicating the sum of sources over a network," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 4, pp. 655–665, April 2013.

[27] N. Silberstein, A. S. Rawat, and S. Vishwanath, "Error resilience in distributed storage via rank-metric codes," in *Proc. 50th Annual Allerton Conf. on Communication, Control, and Computing (Allerton)*, Urbana-Champaign, IL, Oct. 2012, pp. 1150 –1157.

[28] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath, "Optimal locally repairable and secure codes for distributed storage systems," *IEEE Trans. Inf. Theory*, vol. 60, no. 1, pp. 212–236, Jan 2014.

[29] G. M. Kamath, N. Silberstein, N. Prakash, A. S. Rawat, V. Lalitha, O. O. Koyluoglu, P. V. Kumar, and S. Vishwanath, "Explicit mbr all-symbol locality codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2013, pp. 504–508.

[30] Nicolas T Courtois, Matthieu Finiasz, and Nicolas Sendrier, "How to achieve a mceliece-based digital signature scheme," in *Advances in CryptologyASIACRYPT 2001*. Springer, 2001, pp. 157–174.

[31] R. Roth, *Introduction to coding theory*. Cambridge University Press, 2006.

[32] Noga Alon, "Combinatorial Nullstellensatz," *Combinatorics, Probability and Computing*, 1999.

APPENDIX A
PROOF OF LEMMA 2.1

Note that Part 1 is the definition of MRSC given in Definition 1. To see the equivalence of Parts 1 and 2, we simply note that any set $S \subset [n], |S| = k$ is a k-core of $\mathcal{C}_0^{\perp}$ if and only if $\rho(G_0|_S) = k$. Next, to see why Part 2 implies Part 3, let S denote any k-core of $\mathcal{C}_0^{\perp}$, and consider the matrix $H|_{[n]\setminus S}$. If $\rho\left(H|_{[n]\setminus S}\right) < n-k$, then since $\rho(H) = n-k$, this means $\rho\left(H|_S\right) > 0$, which contradicts the assumption in Part 2 that S is a k-core of $\mathcal{C}^{\perp}$ as well. Hence we conclude that Part 2 implies $\rho\left(H|_{[n]\setminus S}\right) = n-k$. The fact that Part 3 implies Part 2 also follows similarly. Next we will prove the equivalence of Parts 1 and 4. Clearly, Part 4 implies Part 1, since in Part 4 we consider all sets of cardinality less than or equal to k. To see why Part 1 implies Part 4, consider any set $S \subset [n], |S| = k-x$, such that $\rho(G_0|_S) = k-x$, for some $x, 1 \leq x \leq k-1$. Since $\rho(G_0) = t > k$, there exists a set $S' \subset [n], |S'| = k-x$ such that $|S \cap S'| = 0$ and $\rho(G_0|_{S\cup S'}) = k$. Now, from Part 1, we know that $\rho(G|_{S\cup S'}) = k$, which then implies that $\rho(G|_S) = k-x$. This completes the proof of Lemma 2.1.

APPENDIX B
PROOF OF THEOREM 2.4

The following fact will be used in the proof of Theorem 2.4.

Fact 1 (Problem 3.33, [31]): Consider the finite field $\mathbb{F}_Q = \mathbb{F}_{q^t}$ as a vector space over F_q, and let $\{\alpha_i \in \mathbb{F}_Q, 1 \leq i \leq k\}$ denote a set of $k, k \leq t$ linearly independent elements over $\mathbb{F}_q$. Then the $k \times k$ matrix

$$\begin{bmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_k \\ \alpha_1^q & \alpha_2^q & \cdots & \alpha_k^q \\ & & \vdots & \\ \alpha_1^{q^{k-1}} & \alpha_2^{q^{k-1}} & \cdots & \alpha_k^{q^{k-1}} \end{bmatrix}. \tag{71}$$

is invertible over $\mathbb{F}_Q$.

The above statement is equivalent to saying that a degree q^{k-1}- linearized polynomial $f(x) \in \mathbb{F}_Q[x]$ given by

$$f(x) \;=\; \sum_{i=0}^{k-1} a_i x^{q^{i-1}},\ a_i \in \mathbb{F}_{q^t} \tag{72}$$

can be uniquely identified given its evaluations at k points $\{\alpha_1, \cdots, \alpha_k\}$ that are linearly independent over $\mathbb{F}_q$.

Proof of Theorem 2.4: Without loss of generality, suppose that $\rho\left(G_0|_{[k]}\right) = k$. It then follows that the elements $\{\beta_1, \cdots, \beta_k\}$ given by

$$[\beta_1 \ \cdots \ \beta_k] \;=\; [\alpha_1 \ \cdots \ \alpha_t] G_0|_{[k]}, \tag{73}$$

are linearly independent over $\mathbb{F}_q$. Combining the above statement along with Fact 1, we get that the matrix $G|_{[k]}$ given by

$$G|_{[k]} \;=\; \begin{bmatrix} \beta_1 & \beta_2 & \cdots & \beta_k \\ \beta_1^q & \beta_2^q & \cdots & \beta_k^q \\ & & \vdots & \\ \beta_1^{q^{k-1}} & \beta_2^{q^{k-1}} & \cdots & \beta_k^{q^{k-1}} \end{bmatrix}. \tag{74}$$

is invertible. The proof now follows from Definition 1 of MRSCs.

APPENDIX C
PROOF OF CONVERSE PART OF LEMMA 4.1

Consider the $[n, t]$ code $\mathcal{C}_0$, and let $\mathcal{C}$ denote an $[n, k]$ MRSC of $\mathcal{C}_0$, where $k = t - \Delta$ for some $\Delta < t$. We will show that there exists an $[n+\Delta, t]$ extension $\mathcal{C}_0^{(e)}$ of $\mathcal{C}_0$ such that for any $S \subseteq [n], |S| = t-\Delta$, we have

$$\rho\left(G_0^{(e)}|_{S\cup\{n+1,\ldots,n+\Delta\}}\right) \;=\; \rho\left(G_0|_S\right) + \Delta. \tag{75}$$

Here, G_0 and $G_0^{(e)}$ denote generator matrices of the codes $\mathcal{C}_0$ and $\mathcal{C}_0^{(e)}$, respectively. Towards this, consider the $(n-k)\times n$ parity check matrix H of $\mathcal{C}$ given by

$$H \;=\; \left[\begin{array}{c} H_0 \\ H_e \end{array}\right], \tag{76}$$

where H_0 is a parity check matrix of $\mathcal{C}$, and H_e is some $(t-k)\times n$ matrix over $\mathbb{F}_q$ such that $\rho(H_e) = t-k$. We define the candidate code $\mathcal{C}_0^{(e)}$ for the extension as one whose parity check matrix $H_0^{(e)}$ is given by

$$H_0^{(e)} \;=\; \left[\begin{array}{cc} H_0 & 0_{n-t\times\Delta} \\ H_e & I_\Delta \end{array}\right]. \tag{77}$$

First of all, to see that $\mathcal{C}_0^{(e)}$ is indeed an extension of $\mathcal{C}_0$, consider the code $\mathcal{C}'$ generated by $G_0^{(e)}|_{[n]}$, and observe that

$$\mathcal{C}'^{\perp} \;=\; \mathcal{C}_0^{\perp}. \tag{78}$$

The above statement follows from the fact that the dual of the punctured code $\mathcal{C}'$ is the shortened code of the dual code. Clearly, (78) implies that $\mathcal{C}'$ = $\mathcal{C}$, and hence $\mathcal{C}_0^{(e)}$ is an extension of $\mathcal{C}_0$. Next, we will show that $\mathcal{C}_0^{(e)}$ satisfies (75). Towards this, consider any set $\mathcal{S}\subset[n], |S| = k' \leq k$ such that $\rho\left(G_0|_S\right) = k'$. We will prove that

$$\rho\left(G_0|_{S\cup\{n+1,\ldots,n+\Delta\}}\right) = k'+\Delta. \tag{79}$$

Note that proving (79) is sufficient to prove (75). We will prove (79) by the method of contradiction by supposing that $\rho\left(G_0|_{S\cup\{n+1,\ldots,n+\Delta\}}\right) < k'+\Delta$. In this case, we know that there exists a non-zero vector $\mathbf{c^{(e)}} \in \left(\mathcal{C}_{\mathbf{0}}^{(\mathbf{e})}\right)^{\perp}$ having a support T such that

- $T \subset S\cup\{n+1,\ldots,n+\Delta\}$,
- $T\cap\{n+1,\ldots,n+\Delta\}$ is not empty, and
- $T\cap S$ is also not empty (this follows since H has full row rank).

Now, if we consider the vector $\mathbf{c} = \mathbf{c^{(e)}}|_{\mathbf{S}}$, this would mean that the non-zero vector $\mathbf{c}\in\mathcal{C}^{\perp}$ (but clearly $\mathbf{c}\notin\mathcal{C}_{\mathbf{0}}^{\perp}$). This then means that $\rho\left(G|_S\right) < k'$. Using Part 4. of Lemma 2.1, we see that this contradicts our assumption that $\mathcal{C}$ is an $[n,k]$ MRSC of $\mathcal{C}_0$. Hence we conclude that (79) is indeed true, and this completes the proof of the converse.

Appendix D
Proof of Lemma 5.2

Suppose that we are given an $[n,t]$ code $\mathcal{C}_0$ over $\mathbb{F}_q$, and an $[n,s]$ subcode $\widehat{\mathcal{C}}$ of $\mathcal{C}_0$ such that the following condition is satisfied:

$$\rho(G_0|_S) = k \quad\Longrightarrow\quad \rho\left(\widehat{G}|_S\right) = s,\ S\subseteq[n], |S| = k, \tag{80}$$

where k is such that $s<k<t$. We need to show that there exists an $[n,k]$ code $\mathcal{C}$ such that

(P1) $\widehat{\mathcal{C}}\subset\mathcal{C}$ and

(P2) $\mathcal{C}$ is an MRSC of $\mathcal{C}_0$,

whenever $q > ((t-k)s+1)\binom{n-1}{k}$. We use the equivalent definition of MRSCs given in Part 3. of Lemma 2.1 for proving the above statement. Let H_0 denote a parity check matrix of the code $\mathcal{C}_0$. A parity check matrix H for the code $\mathcal{C}$ can then be given by $H = \left[\begin{array}{c} H_0 \\ \hline \\ \Delta H \end{array}\right]$, where ΔH is a $(t-k)\times n$ matrix over $\mathbb{F}_q$. We will show that it is possible to identify ΔH under sufficiently large q such that both (P1) and (P2) are satisfied. Towards this, assume that the entries of ΔH are filled with the indeterminates $x_{i,j}, 1\leq i\leq t-k, 1\leq j\leq n$. Also, without loss of generality, assume that the generator matrix of the code $\widehat{\mathcal{C}}$ is given by $\widehat{G} = \left[I_s \mid P_{s\times(n-s)}\right]$. Now, note that in order to satisfy (P1), it is sufficient if the $\{x_{i,j}\}$s satisfy the following relations:

$$\left[\begin{array}{c} x_{i,1} \\ x_{i,2} \\ \vdots \\ x_{i,s} \end{array}\right] \;=\; -P_{s\times(n-s)}\left[\begin{array}{c} x_{i,s+1} \\ x_{i,s+2} \\ \vdots \\ x_{i,n} \end{array}\right],\ \forall\ i\in\{1,\ \cdots\ ,\ t-k\}. \tag{81}$$

The above equation simply says that each of the $(t-k)$ n-length vectors $[x_{i,1},\ \ldots,\ x_{i,n}]^T, 1 \leq i \leq t-k$ must be picked from within the code $\widehat{\mathcal{C}}^{\perp}$. Thus, let us assume that the $\{x_{i,j}\}$s are related as given by (81). Next, in order to satisfy (P2), we know from Part 3 of Lemma 2.1 that it is sufficient if the following property is satisfied:

(P3) for any $S \subseteq [n], |S| = k$ which is a k-core of $\mathcal{C}_0^{\perp}$, we have rank$(H|_{[n]\setminus S}) = n-k$.

Let us now see why it is always possible to pick the $\{x_{i,j}\}$s simultaneously satisfying (81) and (P3). Towards this, let S denote a k-core of $\mathcal{C}_0^{\perp}$, and consider the $(n-k)\times(n-k)$ matrix $H|_{[n]\setminus S}$. For satisfying (P3), we need to pick the $\{x_{i,j}\}$s such that $H|_{[n]\setminus S}$ is invertible. Of course, this can be done only if the matrix $H|_{[n]\setminus S}$ is not trivially rank deficient even before we pick the $\{x_{i,j}\}$s. We claim that there is no such trivial rank deficiency either due to the rows of the matrix H_0 or due to the relations in (81). To see why this is the case, observe the following two points:

- Since S is a k-core of $\mathcal{C}_0^{\perp}$, we have $\rho\left(H_0|_{[n]\setminus S}\right) = n-t$.
- Secondly, if we restrict our attention to $\Delta H|_{[n]\setminus S}$, then the condition in (80) ensures that there is no linear dependency among the columns of $\Delta H|_{[n]\setminus S}$, even when the $\{x_{i,j}\}s$ are picked such that (81) is satisfied. To see why this is true, suppose that the relations in (81) imply a linear dependency among the columns of $\Delta H|_{[n]\setminus S}$. This is possible only if there exists a non-zero vector $\widehat{\mathbf{c}} \in \widehat{\mathcal{C}}$ such that supp$(\widehat{\mathbf{c}}) \subseteq [n]\setminus S$, which then implies that $\rho\left(\widehat{G}|_S\right) < s$. However this contradicts our assumption in (80) that $\rho\left(\widehat{G}|_S\right) = s$, and we thus conclude that the relations in (81) do not introduce any linear dependency among the columns of $\Delta H|_{[n]\setminus S}$.

For the rest of the proof, let us define the multivariate polynomial $f_S(\{x_{i,j}\})$ as $f_S(\{x_{i,j}\}) = \det(H|_{[n]\setminus S})$, and also let $f(\{x_{i,j}\}) = \Pi_{S\in\mathcal{T}} f_S(\{x_{i,j}\})$, where $\mathcal{T}$ denotes the collection of all k-cores of $\mathcal{C}_0^{\perp}$. We substitute for the $\{x_{i,j}, 1 \leq i \leq t-k, 1 \leq j \leq s\}$s in the polynomial $f(\{x_{i,j}\})$ using (81), and then note that the degree of any $x_{i,j}, 1 \leq i \leq t-k, s+1 \leq j \leq n$ in the resultant polynomial is upper bounded by $\binom{n}{k}$. To see why this is true, we make the following two observations:

- The degree of any $x_{i,j}, 1 \leq i \leq t-k, s+1 \leq j \leq n$ in any of the individual polynomials $f_S(\{x_{i,j}\})$ is at most 1. This is true even after we substitute for the $\{x_{i,j}, 1 \leq i \leq t-k, 1 \leq j \leq s\}$s in the polynomial $f_S(\{x_{i,j}\})$ using (81). The reason why the relations in (81) do not increase the degree of any individual $x_{i,j}$ is because of the fact that these relations only relate the elements of the same row in the matrix ΔH, whereas any monomial in the determinant polynomial $f_S(\{x_{i,j}\})$ consists of $n-k$ terms that are picked from $n-k$ distinct rows.
- The cardinality of the set $\mathcal{T}$ is upper bounded by $\binom{n}{k}$.

The proof now follows by an application of Combinatorial Nullstellensatz [32]) which says that there exists a non-zero evaluation of $f(\{x_{i,j}\})$ in $\mathbb{F}_q$, if $q > \binom{n}{k}$.

APPENDIX E
PROOF OF THEOREM 5.4

Without loss of generality, suppose that $\rho\left(G_0|_S\right) = k$, where $S = \{1,\ldots,k\}$. We know from (37) that $\rho\left(\widehat{G}|_S\right) = s$. Once again, without loss of generality, assume that the matrix $\widehat{G}|_S$ is given by

$$\widehat{G}|_S \ = \ \left[I_s \mid \widehat{P}\right], \tag{82}$$

for some $s\times(k-s)$ matrix $\widehat{P}$ over $\mathbb{F}_q$. Consider the matrix B introduced in (34), and assume that $B|_S = [B_1 \ \ B_2]$, where B_1 and B_2 are matrices over $\mathbb{F}_q$ of sizes $(t-s)\times s$ and $(t-s)\times(k-s)$, respectively. The matrix $G_0|_S$ can now be written in terms of these sub matrices as follows:

$$G_0|_S \ = \ \left[\begin{array}{c|c} I_s & \widehat{P} \\ \hline B_1 & B_2 \end{array}\right], \tag{83}$$

which can be row-reduced and written in the following form:

$$\text{row-reduced}\left(G_0|_S\right) \ = \ \left[\begin{array}{c|c} I_s & \widehat{P} \\ \hline 0_{(t-s)\times s} & B_2 - B_1\widehat{P} \end{array}\right]. \tag{84}$$

We now make the following observations given the row-reduced form of $G_0|_S$ in (84).

- Since $\rho\left(G_0|_S\right) = k$, it must be true that

$$\rho\left(B_2 - B_1\widehat{P}\right) = k - s. \tag{85}$$

- Using the fact that $[\alpha_1, \cdots, \alpha_{t-s}]B = [\beta_1, \cdots, \beta_n]$ from (35), we get that

$$\begin{aligned}[\alpha_1, \cdots, \alpha_{t-s}]\left(B_2 - B_1\widehat{P}\right) &= [\beta_{s+1}, \cdots, \beta_k] - [\beta_1, \cdots, \beta_s]\widehat{P} \\ &\triangleq [\gamma_1, \cdots, \gamma_{k-s}],\end{aligned} \tag{86}$$

  where we define the parameter γ_i as follows:

$$\gamma_i = \beta_{s+i} - \sum_{j=1}^{s} \beta_j \widehat{P}(j,i), \quad 1 \leq i \leq k-s. \tag{87}$$

  Here $\widehat{P}(j,i)$ denotes the $(j,i)^{\text{th}}$ entry in the matrix $\widehat{P}$. Since $\widehat{P}$ is a matrix over $\mathbb{F}_q$, we further note that

$$\gamma_i^q = \beta_{s+i}^q - \sum_{j=1}^{s} \beta_j^q \widehat{P}(j,i), \quad 1 \leq i \leq k-s. \tag{88}$$

Let us next consider the matrix $G|_S$ and show that $\rho\left(G|_S\right) = k$. Towards this, first of all note from (36) and (82) that $G|_S$ can be written as follows:

$$G|_S = \left[\begin{array}{c|c} I_s & \widehat{P} \\ \hline \Lambda_1 & \Lambda_2 \end{array}\right], \tag{89}$$

where

$$\Lambda_1 = \begin{bmatrix} \beta_1 & \beta_2 & \cdots & \beta_s \\ \beta_1^q & \beta_2^q & \cdots & \beta_s^q \\ & & \vdots & \\ \beta_1^{q^{k-s-1}} & \beta_2^{q^{k-s-1}} & \cdots & \beta_s^{q^{k-s-1}} \end{bmatrix} \text{ and } \Lambda_2 = \begin{bmatrix} \beta_{s+1} & \beta_{s+2} & \cdots & \beta_k \\ \beta_{s+1}^q & \beta_{s+2}^q & \cdots & \beta_k^q \\ & & \vdots & \\ \beta_{s+1}^{q^{k-s-1}} & \beta_{s+2}^{q^{k-s-1}} & \cdots & \beta_k^{q^{k-s-1}} \end{bmatrix}. \tag{90}$$

The matrix $G|_S$ in (89) can be row-reduced and written as follows:

$$\text{row-reduced}\left(G|_S\right) = \left[\begin{array}{c|c} I_s & \widehat{P} \\ \hline 0_{(k-s)\times s} & \Lambda_2 - \Lambda_1\widehat{P} \end{array}\right]. \tag{91}$$

In order to prove that $\rho\left(G|_S\right) = k$, note that it is enough if we show that $\rho\left(\Lambda_2 - \Lambda_1\widehat{P}\right) = k - s$. Towards this, observe that

$$\begin{aligned}\Lambda_2 - \Lambda_1\widehat{P} &= \begin{bmatrix} [\beta_{s+1}, \cdots, \beta_k] - [\beta_1, \cdots, \beta_s]\,\widehat{P} \\ \left[\beta_{s+1}^q, \cdots, \beta_k^q\right] - \left[\beta_1^q, \cdots, \beta_s^q\right]\widehat{P} \\ \vdots \\ \left[\beta_{s+1}^{q^{k-s-1}}, \cdots, \beta_k^{q^{k-s-1}}\right] - \left[\beta_1^{q^{k-s-1}}, \cdots, \beta_s^{q^{k-s-1}}\right]\widehat{P} \end{bmatrix} \\ &= \begin{bmatrix} \gamma_1 & \gamma_2 & \cdots & \gamma_{k-s} \\ \gamma_1^q & \gamma_2^q & \cdots & \gamma_{k-s}^q \\ & & \vdots & \\ \gamma_1^{q^{k-s-1}} & \gamma_2^{q^{k-s-1}} & \cdots & \gamma_{k-s}^{q^{k-s-1}} \end{bmatrix},\end{aligned} \tag{92}$$

where (92) follows from (88). Recall that the $\{\gamma\}$s are related to the $\{\alpha\}s$ as given in (86). From (85) we know that $\rho\left(B_2 - B_1\widehat{P}\right) = k - s$, from which it follows that the $\{\gamma\}$s are linearly independent over $\mathbb{F}_q$. Now, we can apply Fact 1 to see that the matrix $\Lambda_2 - \Lambda_1\widehat{P}$ in (92) in indeed invertible, i.e., $\rho\left(\Lambda_2 - \Lambda_1\widehat{P}\right) = k - s$. From this it follows that $\rho\left(G|_S\right) = k$, and this completes our proof.